

**QNX[®] Neutrino[®] Realtime
Operating System**

Library Reference

For QNX[®] Neutrino[®] 6.3

QNX Software Systems Ltd.
175 Terence Matthews Crescent
Kanata, Ontario
K2M 1W8
Canada
Voice: +1 613 591-0931
Fax: +1 613 591-3579
Email: info@qnx.com
Web: <http://www.qnx.com/>

© QNX Software Systems Ltd. 2004. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise without the prior written permission of QNX Software Systems Ltd.

Although every precaution has been taken in the preparation of this book, we assume no responsibility for any errors or omissions, nor do we assume liability for damages resulting from the use of the information contained in this book.

Third-party copyright notices

All appropriate copyright notices for third-party software are published in this manual in an appendix called "Third-Party Copyright Notices."

Technical support options

To obtain technical support for any QNX product, visit the **Technical Support** section in the **Support** area on our website (www.qnx.com). You'll find a wide range of support options, including our free web-based QNX Developer's Network.

QNX, Momentics, Neutrino, and Photon are registered trademarks of QNX Software Systems Ltd.
All other trademarks and registered trademarks belong to their respective owners.

Contents

About This Reference xlix

What's new in QNX Neutrino 6.3.0	li
New content	li
Changed content	liii
Errata	liiii
What's new in QNX Neutrino 6.2.1	liii
New content	liiii
Changed content	liiii
Errata	liv
What's new in QNX Neutrino 6.2	lvi
New Content	lvi
Deprecated Content	lix
Errata	lix
What's new in the QNX Neutrino 6.1.0 docs	lix
New content	lix
Deprecated content	lxi

Summary of Functions 1

Summary of function categories	3
Asynchronous I/O functions	6
Atomic functions	7
Character manipulation functions	8
Conversion functions	9
Directory functions	12
Dispatch interface functions	13

File manipulation functions	17
IPC functions	19
Hardware functions	25
Math functions	26
Memory allocation functions	33
Memory manipulation functions	34
Message queue functions	36
Multibyte character functions	37
QNX Neutrino-specific IPC functions	37
Operating system I/O functions	40
PC Card functions	44
Platform-specific functions	44
Process environment functions	46
Process manipulation functions	49
Realtime timer functions	58
Resource manager functions	60
Searching and sorting functions	65
Shared memory functions	67
Signal functions	67
Stream I/O functions	69
String manipulation functions	72
System database functions	75
System message log functions	76
TCP/IP functions	77
Terminal control functions	84
Thread functions	85
Time functions	95
Variable-length argument list functions	97
Wide-character functions	97
What's in a function description?	102
Synopsis:	102
Arguments:	102

Library: 102
Description: 102
Returns: 102
Errors: 103
See also: 103
Examples: 103
Classification: 103
Function safety: 107

Manifests 109

abort() 113
abs() 115
accept() 117
access() 120
acos(), acosf() 123
acosh(), acoshf() 125
addrinfo 127
aio_cancel() 129
aio_error() 131
aio_fsync() 133
aio_read() 135
aio_return() 136
aio_suspend() 138
aio_write() 140
alarm() 141
alloca() 144
alphasort() 147
_amblksiz 149
_argc 150
_argv 151
asctime(), asctime_r() 152
asin(), asinf() 154
asinh(), asinhf() 156

<i>assert()</i>	158
<i>atan()</i> , <i>atanf()</i>	161
<i>atan2()</i> , <i>atan2f()</i>	163
<i>atanh()</i> , <i>atanhf()</i>	165
<i>atexit()</i>	167
<i>atof()</i>	170
<i>atoh()</i>	172
<i>atoi()</i>	174
<i>atol()</i> , <i>atoll()</i>	176
<i>atomic_add()</i>	178
<i>atomic_add_value()</i>	180
<i>atomic_clr()</i>	182
<i>atomic_clr_value()</i>	184
<i>atomic_set()</i>	186
<i>atomic_set_value()</i>	188
<i>atomic_sub()</i>	190
<i>atomic_sub_value()</i>	192
<i>atomic_toggle()</i>	194
<i>atomic_toggle_value()</i>	196
<i>_auxv</i>	198
<i>basename()</i>	199
<i>bcmp()</i>	202
<i>bcopy()</i>	204
<i>bind()</i>	206
<i>bindresvport()</i>	209
<i>brk()</i>	211
<i>bsearch()</i>	214
<i>_btext</i>	217
<i>btowc()</i>	218
<i>bzero()</i>	220
<i>cabs()</i> , <i>cabsf()</i>	222
<i>calloc()</i>	224

<i>cbrt()</i> , <i>cbrtf()</i>	226
<i>ceil()</i> , <i>ceilf()</i>	228
<i>cfmakeraw()</i>	230
<i>cfgetispeed()</i>	232
<i>cfgetospeed()</i>	234
<i>cfgopen()</i>	236
<i>cfree()</i>	240
<i>cfsetispeed()</i>	242
<i>cfsetospeed()</i>	245
<i>ChannelCreate()</i> , <i>ChannelCreate_r()</i>	248
<i>ChannelDestroy()</i> , <i>ChannelDestroy_r()</i>	255
<i>chdir()</i>	258
<i>chmod()</i>	261
<i>chown()</i>	265
<i>chroot()</i>	268
<i>chsize()</i>	271
<i>clearenv()</i>	274
<i>clearerr()</i>	277
<i>clock()</i>	279
<i>ClockAdjust()</i> , <i>ClockAdjust_r()</i>	281
<i>ClockCycles()</i>	284
<i>clock_getcpuclockid()</i>	286
<i>clock_getres()</i>	288
<i>clock_gettime()</i>	290
<i>clock_nanosleep()</i>	293
<i>clock_settime()</i>	297
<i>ClockId()</i> , <i>ClockId_r()</i>	300
<i>ClockPeriod()</i> , <i>ClockPeriod_r()</i>	303
<i>ClockTime()</i> , <i>ClockTime_r()</i>	306
<i>close()</i>	309
<i>closedir()</i>	311
<i>closelog()</i>	314

<i>_cmdfd()</i>	315
<i>_cmdname()</i>	316
<i>confstr()</i>	318
<i>connect()</i>	323
<i>ConnectAttach()</i> , <i>ConnectAttach_r()</i>	326
<i>ConnectClientInfo()</i> , <i>ConnectClientInfo_r()</i>	331
<i>ConnectDetach()</i> , <i>ConnectDetach_r()</i>	335
<i>ConnectFlags()</i> , <i>ConnectFlags_r()</i>	337
<i>ConnectServerInfo()</i> , <i>ConnectServerInfo_r()</i>	340
<i>copysign()</i> , <i>copysignf()</i>	343
<i>cos()</i> , <i>cosf()</i>	345
<i>cosh()</i> , <i>coshf()</i>	347
<i>creat()</i> , <i>creat64()</i>	349
<i>crypt()</i>	353
<i>ctermid()</i>	355
<i>ctime()</i> , <i>ctime_r()</i>	357
<i>daemon()</i>	360
<i>daylight</i>	362
<i>DebugBreak()</i>	363
<i>DebugKDBreak()</i>	365
<i>DebugKDOutput()</i>	366
<i>delay()</i>	368
<i>devctl()</i>	370
<i>difftime()</i>	380
<i>direntl()</i>	382
<i>dirname()</i>	385
<i>dispatch_block()</i>	388
<i>dispatch_context_alloc()</i>	391
<i>dispatch_context_free()</i>	394
<i>dispatch_create()</i>	396
<i>dispatch_destroy()</i>	399
<i>dispatch_handler()</i>	401

<i>dispatch_timeout()</i>	404
<i>dispatch_unblock()</i>	406
<i>div()</i>	408
<i>dladdr()</i>	410
<i>dlclose()</i>	413
<i>dlderror()</i>	415
<i>dlopen()</i>	417
<i>dlsym()</i>	424
<i>dn_comp()</i>	427
<i>dn_expand()</i>	429
<i>drand48()</i>	431
<i>drem(), dremf()</i>	433
<i>ds_clear()</i>	435
<i>ds_create()</i>	437
<i>ds_deregister()</i>	440
<i>ds_flags()</i>	442
<i>ds_get()</i>	444
<i>ds_register()</i>	446
<i>ds_set()</i>	448
<i>dup()</i>	450
<i>dup2()</i>	453
<i>eaccess()</i>	456
<i>_edata</i>	459
<i>encrypt()</i>	460
<i>_end</i>	462
<i>endgrent()</i>	463
<i>endhostent()</i>	464
<i>ENDIAN_BE16()</i>	465
<i>ENDIAN_BE32()</i>	467
<i>ENDIAN_BE64()</i>	469
<i>ENDIAN_LE16()</i>	471
<i>ENDIAN_LE32()</i>	473

<i>ENDIAN_LE64()</i>	475
<i>ENDIAN_RET16()</i>	477
<i>ENDIAN_RET32()</i>	479
<i>ENDIAN_RET64()</i>	481
<i>ENDIAN_SWAP16()</i>	483
<i>ENDIAN_SWAP32()</i>	485
<i>ENDIAN_SWAP64()</i>	487
<i>endnetent()</i>	489
<i>endprotoent()</i>	490
<i>endpwent()</i>	491
<i>endservent()</i>	492
<i>endspent()</i>	493
<i>endutent()</i>	494
<i>environ</i>	495
<i>eof()</i>	496
<i>erand48()</i>	498
<i>erf(), erf()</i>	500
<i>erfc(), erfc()</i>	502
<i>err(), errx()</i>	504
<i>errno</i>	507
<i>_etext</i>	515
<i>execl()</i>	516
<i>execle()</i>	522
<i>execlp()</i>	529
<i>execlpe()</i>	535
<i>execv()</i>	540
<i>execve()</i>	546
<i>execvp()</i>	552
<i>execvpe()</i>	558
<i>_exit()</i>	563
<i>exit()</i>	566
<i>exp(), expf()</i>	569

expm1(), *expm1f()* 571
fabs(), *fabsf()* 574
fcfgopen() 576
fchmod() 578
fchown() 581
fclose() 584
fcloseall() 586
fcntl() 588
fdatasync() 597
fdopen() 599
feof() 602
ferror() 604
fflush() 606
ffs() 608
fgetc() 609
fgetchar() 611
fgetpos() 613
fgets() 615
fgetspent() 618
fgetwc() 621
fgetws() 623
fileno() 626
finite(), *finitef()* 629
fink() 631
flock() 634
flockfile() 637
floor(), *floorf()* 639
flushall() 641
fmod(), *fmodf()* 643
fnmatch() 646
fopen(), *fopen64()* 650
fork() 655

<i>forkpty()</i>	659
<i>fp_exception_mask()</i>	661
<i>fp_exception_value()</i>	664
<i>fp_precision()</i>	667
<i>fp_rounding()</i>	670
<i>fpathconf()</i>	673
<i>fprintf()</i>	676
<i>fputc()</i>	678
<i>fputchar()</i>	680
<i>fputs()</i>	682
<i>fputwc()</i>	684
<i>fputws()</i>	686
<i>fread()</i>	688
<i>free()</i>	691
<i>freeaddrinfo()</i>	693
<i>freeifaddrs()</i>	695
<i>freopen(), freopen64()</i>	697
<i>frexp(), frexpf()</i>	701
<i>fscanf()</i>	703
<i>fseek(), fseeko()</i>	705
<i>fsetpos()</i>	708
<i>fstat(), fstat64()</i>	710
<i>fstatvfs(), fstatvfs64()</i>	714
<i>fsync()</i>	718
<i>ftell(), ftello()</i>	720
<i>ftime()</i>	723
<i>ftruncate(), ftruncate64()</i>	726
<i>ftrylockfile()</i>	729
<i>ftw(), ftw64()</i>	731
<i>funlockfile()</i>	734
<i>futime()</i>	736
<i>fwide()</i>	739

fwprintf() 741
fwrite() 743
fwscanf() 746
gai_strerror() 748
gamma(), *gamma_r()*, *gammaf()*, *gammaf_r()* 750
getaddrinfo() 753
getc() 760
getc_unlocked() 762
getchar() 764
getchar_unlocked() 766
getcwd() 768
getdomainname() 771
getdtablesize() 773
getegid() 775
getenv() 777
geteuid() 779
getgid() 781
getgrent() 783
getgrgid() 786
getgrgid_r() 788
getgrnam() 791
getgrnam_r() 793
getgrouplist() 796
getgroups() 798
gethostbyaddr() 800
gethostbyaddr_r() 803
gethostbyname(), *gethostbyname2()* 806
gethostbyname_r() 809
gethostent() 812
gethostent_r() 814
gethostname() 817
getifaddrs() 819

<i>GETIOVBASE()</i>	821
<i>GETIOVLEN()</i>	823
<i>getitimer()</i>	825
<i>getlogin()</i>	827
<i>getlogin_r()</i>	829
<i>getnameinfo()</i>	831
<i>getnetbyaddr()</i>	836
<i>getnetbyname()</i>	838
<i>getnetent()</i>	840
<i>getopt()</i>	842
<i>getpass()</i>	848
<i>getpeername()</i>	850
<i>getpgid()</i>	852
<i>getpgrp()</i>	854
<i>getpid()</i>	856
<i>getppid()</i>	858
<i>getprio()</i>	860
<i>getprotobyname()</i>	862
<i>getprotobynumber()</i>	864
<i>getprotoent()</i>	866
<i>getpwent()</i>	868
<i>getpwnam()</i>	871
<i>getpwnam_r()</i>	873
<i>getpwuid()</i>	876
<i>getpwuid_r()</i>	878
<i>getrlimit(), getrlimit64()</i>	881
<i>getrusage()</i>	884
<i>gets()</i>	889
<i>getservbyname()</i>	891
<i>getservbyport()</i>	893
<i>getservent()</i>	895
<i>getsid()</i>	897

getsockname() 899
getsockopt() 901
getspent(), *getspent_r()* 911
getspnam(), *getspnam_r()* 915
getsubopt() 918
gettimeofday() 923
getuid() 925
getutent() 927
getutid() 929
getutline() 932
getw() 934
getwc() 936
getwchar() 938
getwd() 940
glob() 942
globfree() 947
gmtime() 949
gmtime_r() 951
h_errno 953
hcreate() 955
hdestroy() 957
herror() 958
hostent 961
hsearch() 963
hsterror() 967
htonl() 969
htons() 971
hwi_find_item() 973
hwi_find_tag() 975
hwi_off2tag() 977
hwi_tag2off() 979
hypot(), *hypotf()* 981

ICMP	983
ICMP6	985
<i>if_freenameindex()</i>	989
<i>if_indextoname()</i>	991
<i>if_nameindex()</i>	993
<i>if_nametoindex()</i>	995
ifaddrs	997
<i>ilogb()</i> , <i>ilogbf()</i>	999
<i>in8()</i>	1001
<i>in8s()</i>	1003
<i>in16()</i> , <i>inbe16()</i> , <i>inle16()</i>	1005
<i>in16s()</i>	1007
<i>in32()</i> , <i>inbe32()</i> , <i>inle32()</i>	1009
<i>in32s()</i>	1011
<i>index()</i>	1013
<i>inet_addr()</i>	1015
<i>inet_aton()</i>	1017
<i>inet_lnaof()</i>	1019
<i>inet_makeaddr()</i>	1021
<i>inet_net_ntop()</i>	1023
<i>inet_netof()</i>	1026
<i>inet_net_pton()</i>	1028
<i>inet_network()</i>	1030
<i>inet_ntoa()</i>	1032
<i>inet_ntoa_r()</i>	1034
<i>inet_ntop()</i>	1036
<i>inet_pton()</i>	1039
<i>inet6_option_*</i> (<i>)</i>	1044
INET6	1051
<i>inet6_rthdr_*</i> (<i>)</i>	1055
<i>initgroups()</i>	1061
<i>initstate()</i>	1063

input_line() 1066
InterruptAttach(), *InterruptAttach_r()* 1068
InterruptAttachEvent(), *InterruptAttachEvent_r()* 1077
InterruptDetach(), *InterruptDetach_r()* 1083
InterruptDisable() 1085
InterruptEnable() 1087
InterruptHookIdle() 1089
InterruptHookTrace() 1093
InterruptLock() 1095
InterruptMask() 1097
InterruptUnlock() 1100
InterruptUnmask() 1102
InterruptWait(), *InterruptWait_r()* 1104
_intr_v86() 1107
_io_connect 1111
_io_connect_fstype_reply 1118
_io_connect_link_reply 1120
ioctl() 1123
iofdinfo() 1125
iofunc_attr_init() 1127
iofunc_attr_lock() 1129
iofunc_attr_t 1131
iofunc_attr_trylock() 1137
iofunc_attr_unlock() 1139
iofunc_check_access() 1141
iofunc_chmod() 1145
iofunc_chmod_default() 1148
iofunc_chown() 1150
iofunc_chown_default() 1153
iofunc_client_info() 1155
iofunc_close_dup() 1157
iofunc_close_dup_default() 1160

<i>iofunc_close_ocb()</i>	1162
<i>iofunc_close_ocb_default()</i>	1164
<i>iofunc_devctl()</i>	1166
<i>iofunc_devctl_default()</i>	1170
<i>iofunc_fdinfo()</i>	1172
<i>iofunc_fdinfo_default()</i>	1175
<i>iofunc_func_init()</i>	1179
<i>iofunc_link()</i>	1182
<i>iofunc_lock()</i>	1186
<i>iofunc_lock_calloc()</i>	1188
<i>iofunc_lock_default()</i>	1190
<i>iofunc_lock_free()</i>	1193
<i>iofunc_lock_ocb_default()</i>	1195
<i>iofunc_lseek()</i>	1197
<i>iofunc_lseek_default()</i>	1200
<i>iofunc_mknod()</i>	1202
<i>iofunc_mmap()</i>	1205
<i>iofunc_mmap_default()</i>	1209
<i>iofunc_notify()</i>	1211
<i>iofunc_notify_remove()</i>	1218
<i>iofunc_notify_trigger()</i>	1220
<i>iofunc_ocb_attach()</i>	1223
<i>iofunc_ocb_calloc()</i>	1225
<i>iofunc_ocb_detach()</i>	1228
<i>iofunc_ocb_free()</i>	1231
iofunc_ocb_t	1233
<i>iofunc_open()</i>	1236
<i>iofunc_open_default()</i>	1241
<i>iofunc_openfd()</i>	1243
<i>iofunc_openfd_default()</i>	1247
<i>iofunc_pathconf()</i>	1249
<i>iofunc_pathconf_default()</i>	1252

iofunc_read_default() 1254
iofunc_read_verify() 1256
iofunc_readlink() 1260
iofunc_rename() 1263
iofunc_space_verify() 1267
iofunc_stat() 1271
iofunc_stat_default() 1273
iofunc_sync() 1276
iofunc_sync_default() 1278
iofunc_sync_verify() 1280
iofunc_time_update() 1283
iofunc_unblock() 1285
iofunc_unblock_default() 1287
iofunc_unlink() 1290
iofunc_unlock_ocb_default() 1293
iofunc_utime() 1295
iofunc_utime_default() 1298
iofunc_write_default() 1301
iofunc_write_verify() 1303
ionotify() 1307
IP 1313
IPsec 1320
ipsec_dump_policy() 1328
ipsec_get_policylen() 1330
ipsec_strerror() 1332
ipsec_set_policy() 1334
IP6 1338
isalnum() 1348
isalpha() 1350
isascii() 1352
isatty() 1354
iscntrl() 1356

isdigit() 1358
isfdtype() 1360
isgraph() 1362
isinf(), *isinff()* 1364
islower() 1366
isnan(), *isnanf()* 1368
isprint() 1370
ispunct() 1372
isspace() 1374
isupper() 1377
iswalnum() 1379
iswalpha() 1381
iswcntrl() 1383
iswctype() 1385
iswdigit() 1387
iswgraph() 1389
iswlower() 1391
iswprint() 1393
iswpunct() 1395
iswspace() 1397
iswupper() 1399
iswxdigit() 1401
isxdigit() 1403
itoa() 1405
j0(), *j0f()* 1408
j1(), *j1f()* 1410
jn(), *jnf()* 1412
jrand48() 1414
kill() 1416
killpg() 1419
labs() 1421
lchown() 1423

lcong48() 1426
ldexp(), *ldexpf()* 1428
ldiv() 1430
lfind() 1432
lgamma(), *lgamma_r()*, *lgammaf()*, *lgammaf_r()* 1435
link() 1438
lio_listio() 1442
listen() 1447
localeconv() 1449
localtime() 1454
localtime_r() 1456
lockf() 1458
log(), *logf()* 1462
log1p(), *log1pf()* 1464
log10(), *log10f()* 1466
logb(), *logbf()* 1468
login_tty() 1471
longjmp() 1473
rand48() 1476
lsearch() 1478
lseek(), *lseek64()* 1481
lstat(), *lstat64()* 1485
ltoa(), *lltoa()* 1488
ltrunc() 1491
main() 1495
mallinfo() 1498
malloc() 1500
mallopt() 1502
max() 1504
mblen() 1506
mbrlen() 1509
mbrtowc() 1511

mbsinit() 1514
mbsrtowcs() 1516
mbstowcs() 1518
mbtowc() 1521
mcheck() 1524
mem_offset(), *mem_offset64()* 1526
memalign() 1530
memccpy() 1532
memchr() 1534
memcmp() 1536
memcpy() 1538
memcpyv() 1540
memicmp() 1542
memmove() 1544
memset() 1546
message_attach() 1548
message_connect() 1555
message_detach() 1558
min() 1561
mkdir() 1563
mkfifo() 1566
mknod() 1569
mkstemp() 1573
mktemp() 1575
mktime() 1577
mlock() 1580
mlockall() 1582
mmap(), *mmap64()* 1584
mmap_device_io() 1591
mmap_device_memory() 1593
modem_open() 1597
modem_read() 1601

modem_script() 1604
modem_write() 1612
modf(), *modff()* 1615
mount() 1617
mount_parse_generic_args() 1620
mprobe() 1623
mprotect() 1625
mq_close() 1628
mq_getattr() 1630
mq_notify() 1633
mq_open() 1636
mq_receive() 1640
mq_send() 1643
mq_setattr() 1646
mq_timedreceive() 1648
mq_timedsend() 1651
mq_unlink() 1654
mrnd48() 1656
_msg_info 1658
MsgDeliverEvent(), *MsgDeliverEvent_r()* 1661
MsgError(), *MsgError_r()* 1669
MsgInfo(), *MsgInfo_r()* 1672
MsgKeyData(), *MsgKeyData_r()* 1674
MsgRead(), *MsgRead_r()* 1682
MsgReadv(), *MsgReadv_r()* 1686
MsgReceive(), *MsgReceive_r()* 1689
MsgReceivePulse(), *MsgReceivePulse_r()* 1694
MsgReceivePulsev(), *MsgReceivePulsev_r()* 1697
MsgReceivev(), *MsgReceivev_r()* 1700
MsgReply(), *MsgReply_r()* 1704
MsgReplyv(), *MsgReplyv_r()* 1707
MsgSend(), *MsgSend_r()* 1710

MsgSendnc(), *MsgSendnc_r()* 1714
MsgSendPulse(), *MsgSendPulse_r()* 1718
MsgSendsv(), *MsgSendsv_r()* 1722
MsgSendsvnc(), *MsgSendsvnc_r()* 1726
MsgSendv(), *MsgSendv_r()* 1730
MsgSendvnc(), *MsgSendvnc_r()* 1734
MsgSendvs(), *MsgSendvs_r()* 1738
MsgSendvsnc(), *MsgSendvsnc_r()* 1742
MsgVerifyEvent(), *MsgVerifyEvent_r()* 1746
MsgWrite(), *MsgWrite_r()* 1748
MsgWritev(), *MsgWritev_r()* 1752
msync() 1755
munlock() 1758
munlockall() 1760
munmap() 1762
munmap_device_io() 1764
munmap_device_memory() 1766
name_attach() 1768
name_close() 1775
name_detach() 1777
name_open() 1779
nanosleep() 1782
nanospin() 1784
nanospin_calibrate() 1786
nanospin_count() 1789
nanospin_ns() 1791
nanospin_ns_to_count() 1793
nap() 1796
napms() 1797
nbaconnect() 1798
nbaconnect_result() 1801
ND_NODE_CMP() 1803

netent 1805
netmgr_ndtostr() 1806
netmgr_remote_nd() 1812
netmgr_strtond() 1814
nextafter(), *nextafterf()* 1816
nftw(), *nftw64()* 1819
nice() 1823
nrnd48() 1825
nsec2timespec() 1827
ntohl() 1829
ntohs() 1831
offsetof() 1833
open(), *open64()* 1835
opendir() 1843
openfd() 1846
openlog() 1849
openpty() 1852
out8() 1854
out8s() 1856
out16(), *outbe16()*, *outle16()* 1858
out16s() 1860
out32(), *outbe32()*, *outle32()* 1862
out32s() 1864
pathconf() 1866
pathfind(), *pathfind_r()* 1870
pathmgr_symlink() 1874
pathmgr_unlink() 1876
pause() 1878
pccard_arm() 1880
pccard_attach() 1884
pccard_detach() 1886
pccard_info() 1888

pccard_lock() 1891
pccard_raw_read() 1893
pccard_unlock() 1895
pci_attach() 1897
pci_attach_device() 1899
pci_detach() 1908
pci_detach_device() 1910
pci_find_class() 1912
pci_find_device() 1914
pci_irq_routing_options() 1916
pci_map_irq() 1919
pci_present() 1921
pci_read_config() 1924
pci_read_config8() 1926
pci_read_config16() 1928
pci_read_config32() 1930
pci_rescan_bus() 1932
pci_write_config() 1934
pci_write_config8() 1937
pci_write_config16() 1939
pci_write_config32() 1941
pclose() 1943
perror() 1945
pipe() 1947
poll() 1949
popen() 1955
posix_mem_offset(), posix_mem_offset64() 1959
posix_memalign() 1961
pow(), powf() 1963
pread(), pread64() 1965
printf() 1968
procmgr_daemon() 1978

procmgr_event_notify() 1980
procmgr_event_trigger() 1985
procmgr_guardian() 1987
procmgr_session() 1990
–progrname 1993
protoent 1994
pthread_abort() 1995
pthread_atfork() 1997
pthread_attr_destroy() 1999
pthread_attr_getdetachstate() 2001
pthread_attr_getguardsize() 2003
pthread_attr_getinheritsched() 2005
pthread_attr_getschedparam() 2007
pthread_attr_getschedpolicy() 2009
pthread_attr_getscope() 2011
pthread_attr_getstackaddr() 2013
pthread_attr_getstacklazy() 2015
pthread_attr_getstacksize() 2017
pthread_attr_init() 2019
pthread_attr_setdetachstate() 2022
pthread_attr_setguardsize() 2024
pthread_attr_setinheritsched() 2027
pthread_attr_setschedparam() 2029
pthread_attr_setschedpolicy() 2031
pthread_attr_setscope() 2033
pthread_attr_setstackaddr() 2035
pthread_attr_setstacklazy() 2037
pthread_attr_setstacksize() 2039
pthread_barrier_destroy() 2041
pthread_barrier_init() 2043
pthread_barrier_wait() 2045
pthread_barrierattr_destroy() 2047

<i>pthread_barrierattr_getpshared()</i>	2049
<i>pthread_barrierattr_init()</i>	2051
<i>pthread_barrierattr_setpshared()</i>	2053
<i>pthread_cancel()</i>	2055
<i>pthread_cleanup_pop()</i>	2057
<i>pthread_cleanup_push()</i>	2059
<i>pthread_cond_broadcast()</i>	2062
<i>pthread_cond_destroy()</i>	2064
<i>pthread_cond_init()</i>	2066
<i>pthread_cond_signal()</i>	2068
<i>pthread_cond_timedwait()</i>	2070
<i>pthread_cond_wait()</i>	2074
<i>pthread_condattr_destroy()</i>	2077
<i>pthread_condattr_getclock()</i>	2079
<i>pthread_condattr_getpshared()</i>	2081
<i>pthread_condattr_init()</i>	2083
<i>pthread_condattr_setclock()</i>	2085
<i>pthread_condattr_setpshared()</i>	2087
<i>pthread_create()</i>	2089
<i>pthread_detach()</i>	2094
<i>pthread_equal()</i>	2096
<i>pthread_exit()</i>	2098
<i>pthread_getconcurrency()</i>	2100
<i>pthread_getcpuclockid()</i>	2102
<i>pthread_getschedparam()</i>	2104
<i>pthread_getspecific()</i>	2106
<i>pthread_join()</i>	2108
<i>pthread_key_create()</i>	2110
<i>pthread_key_delete()</i>	2114
<i>pthread_kill()</i>	2116
<i>pthread_mutex_destroy()</i>	2118
<i>pthread_mutex_getprioceiling()</i>	2120

pthread_mutex_init() 2122
pthread_mutex_lock() 2124
pthread_mutex_setprioceiling() 2128
pthread_mutex_timedlock() 2130
pthread_mutex_trylock() 2133
pthread_mutex_unlock() 2135
pthread_mutexattr_destroy() 2137
pthread_mutexattr_getprioceiling() 2139
pthread_mutexattr_getprotocol() 2141
pthread_mutexattr_getpshared() 2143
pthread_mutexattr_getrecursive() 2145
pthread_mutexattr_gettype() 2147
pthread_mutexattr_init() 2150
pthread_mutexattr_setprioceiling() 2152
pthread_mutexattr_setprotocol() 2154
pthread_mutexattr_setpshared() 2156
pthread_mutexattr_setrecursive() 2158
pthread_mutexattr_settype() 2160
pthread_once() 2163
pthread_rwlock_destroy() 2166
pthread_rwlock_init() 2168
pthread_rwlock_rdlock() 2171
pthread_rwlock_timedrdlock() 2173
pthread_rwlock_timedwrlock() 2176
pthread_rwlock_tryrdlock() 2179
pthread_rwlock_trywrlock() 2181
pthread_rwlock_unlock() 2183
pthread_rwlock_wrlock() 2185
pthread_rwlockattr_destroy() 2187
pthread_rwlockattr_getpshared() 2189
pthread_rwlockattr_init() 2191
pthread_rwlockattr_setpshared() 2193

pthread_self() 2195
pthread_setcancelstate() 2196
pthread_setcanceltype() 2198
pthread_setconcurrency() 2200
pthread_setschedparam() 2202
pthread_setspecific() 2204
pthread_sigmask() 2206
pthread_sleepon_broadcast() 2208
pthread_sleepon_lock() 2210
pthread_sleepon_signal() 2212
pthread_sleepon_timedwait() 2214
pthread_sleepon_unlock() 2218
pthread_sleepon_wait() 2220
pthread_spin_destroy() 2224
pthread_spin_init() 2226
pthread_spin_lock() 2228
pthread_spin_trylock() 2230
pthread_spin_unlock() 2232
pthread_testcancel() 2234
pthread_timedjoin() 2235
_pulse 2238
pulse_attach() 2240
pulse_detach() 2244
putc() 2247
putc_unlocked() 2249
putchar() 2251
putchar_unlocked() 2253
putenv() 2255
puts() 2258
putspent() 2260
pututline() 2263
putw() 2266

putwc() 2268
putwchar() 2270
pwrite(), pwrite64() 2272
qnx_crypt() 2275
qsort() 2277
Raccept() 2281
raise() 2283
rand() 2286
rand_r() 2288
random() 2290
Rbind() 2293
rcmd() 2295
Rconnect() 2298
rdchk() 2300
re_comp() 2302
re_exec() 2304
read() 2306
read_main_config_file() 2311
readblock() 2315
readcond() 2318
readdir() 2324
readdir_r() 2328
readlink() 2331
readv() 2334
realloc() 2338
realpath() 2341
recv() 2343
recvfrom() 2346
recvmsg() 2350
regcomp() 2354
regerror() 2359
regexec() 2361

regfree() 2364
remainder(), remainderf() 2366
remove() 2368
rename() 2371
res_init() 2374
res_mkquery() 2377
res_query() 2380
res_querydomain() 2383
res_search() 2386
res_send() 2389
resmgr_attach() 2392
resmgr_block() 2401
resmgr_connect_funcs_t 2404
resmgr_context_alloc() 2406
resmgr_context_free() 2409
resmgr_context_t 2411
resmgr_detach() 2413
resmgr_devino() 2417
resmgr_handle_tune() 2420
_resmgr_handle_grow() 2423
resmgr_handler() 2425
_resmgr_io_func() 2428
resmgr_io_funcs_t 2430
resmgr_iofuncs() 2436
resmgr_msgread() 2438
resmgr_msgreadv() 2440
resmgr_msgwrite() 2442
resmgr_msgwritev() 2444
_RESMGR_NPARTS() 2446
_resmgr_ocb() 2448
resmgr_open_bind() 2450
resmgr_pathname() 2453

_RESMGR_PTR() 2456
_RESMGR_STATUS() 2458
resmgr_unbind() 2460
rewind() 2462
rewinddir() 2465
Rgetsockname() 2468
rindex() 2470
rint(), rintf() 2472
Rlisten() 2475
rmdir() 2477
ROUTE 2480
Rrcmd() 2488
rresvport() 2490
Rselect() 2492
rsrddbmgr_attach() 2494
rsrddbmgr_create() 2501
rsrddbmgr_destroy() 2505
rsrddbmgr_detach() 2507
rsrddbmgr_devno_attach() 2509
rsrddbmgr_devno_detach() 2513
rsrddbmgr_query() 2515
ruserok() 2518
sbrk() 2520
scalb(), scalbf() 2523
scalbn(), scalbnf() 2526
_scalloc() 2529
scandir() 2531
scanf() 2533
sched_getparam() 2542
sched_get_priority_adjust() 2545
sched_get_priority_max() 2547
sched_get_priority_min() 2549

sched_getscheduler() 2551
sched_param 2553
sched_rr_get_interval() 2559
sched_setparam() 2561
sched_setscheduler() 2564
sched_yield() 2567
SchedGet(), *SchedGet_r()* 2570
SchedInfo(), *SchedInfo_r()* 2573
SchedSet(), *SchedSet_r()* 2576
SchedYield(), *SchedYield_r()* 2579
sctp_bindx() 2581
sctp_connectx() 2584
sctp_freeladdrs() 2586
sctp_freepaddrs() 2587
sctp_getladdrs() 2588
sctp_getpaddrs() 2590
sctp_peeloff() 2592
SCTP 2594
sctp_recvmsg() 2596
sctp_sendmsg() 2598
searchenv() 2602
seed48() 2605
seekdir() 2607
select() 2609
select_attach() 2615
select_detach() 2619
select_query() 2622
sem_close() 2625
sem_destroy() 2627
sem_getvalue() 2629
sem_init() 2631
sem_open() 2634

sem_post() 2639
sem_timedwait() 2641
sem_trywait() 2644
sem_unlink() 2646
sem_wait() 2648
send() 2650
sendmsg() 2653
sendto() 2656
servent 2659
setbuf() 2660
setbuffer() 2662
setdomainname() 2664
setegid() 2666
setenv() 2669
seteuid() 2672
setgid() 2675
setgrent() 2678
setgroups() 2680
sethostent() 2682
sethostname() 2684
SETIOV() 2686
setitimer() 2688
setjmp() 2691
setkey() 2694
setlinebuf() 2696
setlocale() 2698
setlogmask() 2701
setnetent() 2703
setpgid() 2705
setpgrp() 2708
setprio() 2709
setprotoent() 2711

setpwent() 2713
setregid() 2714
setreuid() 2717
setrlimit(), *setrlimit64()* 2719
setservent() 2725
setsid() 2727
setsockopt() 2729
setspent() 2732
setstate() 2733
settimeofday() 2735
setuid() 2737
setutent() 2740
setvbuf() 2742
_sfree() 2745
shm_ctl() 2747
shm_open() 2753
shm_unlink() 2760
shutdown() 2762
sigaction() 2764
sigaddset() 2770
sigblock() 2772
sigdelset() 2774
sigemptyset() 2776
sigevent 2778
sigfillset() 2783
sigismember() 2785
siglongjmp() 2787
sigmask() 2789
signal() 2791
SignalAction(), *SignalAction_r()* 2795
SignalKill(), *SignalKill_r()* 2803
SignalProcmask(), *SignalProcmask_r()* 2809

SignalSuspend(), *SignalSuspend_r()* 2813
SignalWaitinfo(), *SignalWaitinfo_r()* 2816
significand(), *significandf()* 2819
sigpause() 2822
sigpending() 2824
sigprocmask() 2826
sigqueue() 2829
sigsetjmp() 2832
sigsetmask() 2834
sigsuspend() 2836
sigtimedwait() 2838
sigunblock() 2841
sigwait() 2843
sigwaitinfo() 2845
sin(), *sinf()* 2847
sinh(), *sinhf()* 2849
sleep() 2851
_sleepon_broadcast() 2853
_sleepon_destroy() 2855
_sleepon_init() 2857
_sleepon_lock() 2859
_sleepon_signal() 2861
_sleepon_unlock() 2863
_sleepon_wait() 2865
slogb() 2867
slogf() 2869
slogi() 2873
_smalloc() 2875
snmp_close() 2877
snmp_free_pdu() 2879
snmp_open() 2881
snmp_pdu 2883

snmp_pdu_create() 2887
snmp_read() 2889
snmp_select_info() 2891
snmp_send() 2894
snmp_session 2897
snmp_timeout() 2901
snprintf() 2903
socketmark() 2906
socket() 2908
socketpair() 2912
SOCKSinit() 2915
sopen() 2917
sopenfd() 2922
spawn() 2925
spawnl() 2933
spawnle() 2938
spawnlp() 2943
spawnlpe() 2947
spawnp() 2952
spawnv() 2960
spawnve() 2965
spawnvp() 2969
spawnvpe() 2973
sprintf() 2978
sqrt(), *sqrtf()* 2980
srand() 2982
srand48() 2984
srandom() 2986
_srealloc() 2988
sscanf() 2991
stat(), *stat64()* 2993
statvfs(), *statvfs64()* 3001

stderr 3005
stdin 3006
stdout 3007
straddstr() 3008
strcasecmp() 3010
strcat() 3013
strchr() 3015
strcmp() 3017
strcmpi() 3019
strcoll() 3021
strcpy() 3023
strcspn() 3025
strdup() 3027
strerror() 3029
strftime() 3031
stricmp() 3037
strlen() 3039
strlwr() 3041
strncasecmp() 3043
strncat() 3046
strncmp() 3048
strncpy() 3050
strnicmp() 3052
strnset() 3054
strpbrk() 3056
strrchr() 3058
strrev() 3060
strsep() 3062
strset() 3064
strsignal() 3066
strspn() 3068
strstr() 3070

strtod() 3072
strtoimax(), strtoumax() 3075
strtok() 3077
strtok_r() 3080
strtol(), strtoll() 3082
strtoul(), strtoull() 3085
strupr() 3088
strxfrm() 3090
swab() 3093
swprintf() 3095
swscanf() 3097
symlink() 3099
sync() 3102
SyncCondvarSignal(), SyncCondvarSignal_r() 3104
SyncCondvarWait(), SyncCondvarWait_r() 3107
SyncCtl(), SyncCtl_r() 3111
SyncDestroy(), SyncDestroy_r() 3114
SyncMutexEvent(), SyncMutexEvent_r() 3117
SyncMutexLock(), SyncMutexLock_r() 3119
SyncMutexRevive(), SyncMutexRevive_r() 3122
SyncMutexUnlock(), SyncMutexUnlock_r() 3124
SyncSemPost(), SyncSemPost_r() 3127
SyncSemWait(), SyncSemWait_r() 3129
SyncTypeCreate(), SyncTypeCreate_r() 3132
sysconf() 3136
sysctl() 3139
syslog() 3147
sysmgr_reboot() 3150
SYSPAGE_CPU_ENTRY() 3152
SYSPAGE_ENTRY() 3154
_syspage_ptr 3157
system() 3158

tan(), *tanf()* 3161
tanh(), *tanhf()* 3163
tcdrain() 3165
tcdropline() 3167
tcflow() 3170
tcflush() 3173
tcgetattr() 3176
tcgetpgrp() 3178
tcgetsid() 3180
tcgetsize() 3182
tcinject() 3184
tcischars() 3187
TCP 3189
tcsendbreak() 3192
tcsetattr() 3194
tcsetpgrp() 3197
tcsetsid() 3200
tcsetsize() 3202
tell(), *tell64()* 3204
telldir() 3207
tempnam() 3209
termios 3211
thread_pool_control() 3215
thread_pool_create() 3218
thread_pool_destroy() 3225
thread_pool_limits() 3228
thread_pool_start() 3231
ThreadCancel(), *ThreadCancel_r()* 3234
ThreadCreate(), *ThreadCreate_r()* 3238
ThreadCtl(), *ThreadCtl_r()* 3245
ThreadDestroy(), *ThreadDestroy_r()* 3249
ThreadDetach(), *ThreadDetach_r()* 3252

ThreadJoin(), *ThreadJoin_r()* 3254
time() 3257
timer_create() 3259
timer_delete() 3263
timer_getexpstatus() 3265
timer_getoverrun() 3267
timer_gettime() 3269
timer_settime() 3271
timer_timeout(), *timer_timeout_r()* 3274
TimerAlarm(), *TimerAlarm_r()* 3281
TimerCreate(), *TimerCreate_r()* 3284
TimerDestroy(), *TimerDestroy_r()* 3288
TimerInfo(), *TimerInfo_r()* 3290
TimerSettime(), *TimerSettime_r()* 3294
TimerTimeout(), *TimerTimeout_r()* 3298
times() 3306
timespec 3309
timespec2nsec() 3310
timezone 3312
tm 3313
tmpfile(), *tmpfile64()* 3315
tmpnam() 3318
tolower() 3321
toupper() 3323
towctrans() 3325
towlower() 3327
towupper() 3329
TraceEvent() 3331
truncate() 3334
ttyname() 3337
ttyname_r() 3339
tzname 3341

tzset() 3342
ualarm() 3345
UDP 3348
ultoa(), *ulltoa()* 3350
umask() 3353
umount() 3356
UNALIGNED_PUT16() 3358
UNALIGNED_PUT32() 3360
UNALIGNED_PUT64() 3362
UNALIGNED_RET16() 3364
UNALIGNED_RET32() 3366
UNALIGNED_RET64() 3368
uname() 3370
ungetc() 3373
ungetwc() 3375
UNIX 3377
unlink() 3380
unsetenv() 3383
usleep() 3385
utime() 3387
utimes() 3390
utmp 3393
utmpname() 3395
utoa() 3397
va_arg() 3400
va_copy() 3406
va_end() 3408
va_start() 3410
valloc() 3412
verr(), *verrx()* 3414
vfork() 3416
vfprintf() 3418

vfscanf() 3421
vwprintf() 3424
vfwscanf() 3426
vprintf() 3428
vscanf() 3430
vsnprintf() 3433
vsprintf() 3435
vsscanf() 3438
vswprintf() 3441
vswscanf() 3444
vsyslog() 3446
vwarn(), *vwarnx()* 3450
vwprintf() 3452
vwscanf() 3454
wait() 3456
wait3() 3459
wait4() 3462
waitid() 3466
waitpid() 3469
warn(), *warnx()* 3472
wcrtomb() 3474
wscat() 3476
wcschr() 3478
wscmp() 3480
wscoll() 3482
wscopy() 3484
wscspn() 3486
wcsftime() 3488
wcslen() 3490
wcsncat() 3492
wcsncmp() 3494

wcsncpy() 3496
wcsprk() 3498
wcsrchr() 3500
wcsrtombs() 3502
wcsspn() 3504
wcsstr() 3506
westod(), *westof()*, *westold()* 3508
wstoimax(), *wcstoumax()* 3511
wstok() 3513
wstol(), *wcstoll()* 3515
wstombs() 3518
wstoul(), *wcstoull()* 3521
wscxfrm() 3524
wctob() 3526
wctomb() 3528
wctrans() 3531
wctype() 3533
wmemchr() 3535
wmemcmp() 3537
wmemcpy() 3539
wmemmove() 3541
wmemset() 3543
wordexp() 3545
wordfree() 3547
wprintf() 3548
write() 3550
writeblock() 3555
writev() 3558
wscanf() 3561
y0(), *y0f()* 3563
y1(), *y1f()* 3565
yn(), *ynf()* 3567

A SOCKS — A Basic Firewall 3569

About SOCKS 3571

How to SOCKSify a client 3571

What SOCKS expects 3572

B Third-Party Copyright Notices 3575

BSD Stack 3577

BSD Stack and Various Utilities 3578

MINIX Operating System 3585

Regular Expression Handling 3586

Remote Procedure Call (RPC) 3587

SNMPv2 3587

SOCKS 3588

C Summary of Safety Information 3591

Cancellation points 3593

Interrupt handlers 3598

Signal handlers 3601

Multithreaded programs 3614

Glossary 3617

Index 3641

List of Figures

- A hierarchy of processes. 1662
- A deadlock when sending messages improperly among processes.
1663
- MsgSendv()*, client to process manager. 1675
- MsgReplyv()*, process manager to client. 1675
- MsgSendv()*, client to filesystem manager 1676
- Components of a fully qualified pathname. 1807
- Specifying a guardian for child processes. 1987
- Conditions that satisfy an input request. 2319
- Most of the *spawn*()* functions do a lot of work before a message
is sent to **procnto**. 2929



About This Reference



The *Library Reference* describes the C functions, data types, and protocols that are included as part of the QNX Neutrino RTOS.

The *Library Reference* also contains:

- Summary listings of the library, including a description of what you'll find in a function description
- Summary of safety information:
 - functions that are cancellation points
 - functions that you can safely call from an interrupt handler
 - functions that you can safely call from a signal handler
 - functions that you *can't* safely call from a multithreaded program.
- descriptions of manifests
- SOCKS — A Basic Firewall
- Third-Party Copyright Notices
- Glossary

What's new in QNX Neutrino 6.3.0

New content

<i>fopen64()</i>	Large-file support for <i>fopen()</i> .
<i>freopen64()</i>	Large-file support for <i>freopen()</i> .
<i>ftw64()</i>	Large-file support for <i>ftw()</i> .
<i>getnameinfo()</i>	Perform address-to-nodename translation.
<i>inet6_option_*</i> ()	Manipulate IPv6 hop-by-hop and destination options.
<i>inet6_rthdr_*</i> ()	manipulate IPv6 Router header options.

<i>ipsec_dump_policy()</i>	Generate a readable string from an IPsec policy specification.
<i>ipsec_get_policylen()</i>	Get the length of the IPsec policy.
<i>ipsec_set_policy()</i>	Generate an IPsec policy specification structure from a readable string.
<i>nftw()</i> , <i>nftw64()</i>	Walk a file tree and its large-file support.
<i>poll()</i>	Input/output multiplexing.
<i>resmgr_handle_tune()</i>	Tune aspects of client <i>fd</i> -to-OCB mapping
<i>sctp_bindx()</i>	Add or remove one or more addresses from a given association.
<i>sctp_connectx()</i>	Help associate an endpoint that is multi-homed.
<i>sctp_freeladdrs()</i>	Free all resources allocated by <i>sctp_getladdrs()</i> .
<i>sctp_freepaddrs()</i>	Free all resources allocated by <i>sctp_getpaddrs()</i> .
<i>sctp_getladdrs()</i>	Return all locally bound addresses on a socket.
<i>sctp_getpaddrs()</i>	Return all peer addresses in an association.
<i>sctp_peeloff()</i>	Branch off an association into a separate socket.
SCTP	Stream Control Transmission Protocol.
<i>sctp_recvmag()</i>	Receive message using advanced SCTP features.
<i>sctp_sendmag()</i>	Send message using advanced SCTP features.
<i>tmpfile64()</i>	Large-file support for <i>tmpfile()</i> .

Changed content

Errata

What's new in QNX Neutrino 6.2.1

New content

dispatch_unblock()

Unblock all of the threads that are blocked on a dispatch handle

errno

Each thread in a multi-threaded program has its own error value in its thread local storage. No matter which thread you're in, you can simply refer to *errno* — it's defined in such a way that it refers to the correct variable for the thread. For more information, see “Local storage for private data” in the documentation for *ThreadCreate()*.

pthread_attr_setschedpolicy()

Sporadic scheduling (SCHED_SPORADIC) is a new feature of QNX Neutrino 6.2.0.

sched_param

Structure of scheduling parameters

va_copy()

Make a copy of a variable argument list

Changed content

bind(), *bindresvport()*

These functions aren't cancellation points any more, because this conflicted with POSIX.

htonl(), *htons()*, *inet_ntop()*, *inet_pton()*, *isfdtype()*, *ntohl()*, *ntohs()*

These functions have been moved from **libsocket** to **libc**.

Errata

alphasort() This function compares two directory entries; it doesn't sort an array of entries.

execlpe(), *execvpe()*

You can now execute a shell script.

fgetc(), *fgetchar()*, *fgets()*, *fgetwc()*, *fgetws()*, *getc()*, *getc_unlocked()*, *getchar()*, *getchar_unlocked()*, *gets()*, *getw()*, *getwc()*, *getwchar()*

Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

fstat(), *fstat64()* These functions return -1 if an error occurs.

iofunc_mmap(), *iofunc_mmap_default()*

These functions return a nonpositive value on success.

InterruptAttach(), *InterruptAttachEvent()*

You should always set `_NTO_INTR_FLAGS_TRK_MSK`.

mq_getattr(), *mq_setattr()*

The *mq_flags* member of the `mq_attr` structure applies to the message-queue description (i.e. locally), not to the queue as a whole.

mq_open() Corrected the interpretation of the *name* argument.

MsgError(), *MsgError_r()*

If the *error* argument is EOK, the *MsgSend*()* call returns EOK; if *error* is any other value, the *MsgSend*()* call returns -1.

MsgSendPulse(), *MsgSendPulse_r()*

You can now send pulses across the network.

You can send a pulse to any process — not just to a process in the same process group — if your process has the appropriate permission.

name_open() This function returns a nonnegative integer representing a side-channel connection ID, or -1 if an error occurred.

printf() The exponent produced for the **e** and **E** formats is at least two digits long.
Clarified what happens if the format string includes invalid multibyte characters.

pthread_mutex_timedlock(), *pthread_rwlock_timedrdlock()*,
pthread_rwlock_timedwrlock()
The timeout is based on the CLOCK_REALTIME clock.

_resmgr_ocb() Corrected the name.

select() This function and the associated macros are now defined in `<sys/select.h>`, instead of `<sys/time.h>` (which includes `<sys/select.h>`).

sem_open() Corrected the interpretation of the *sem_name* argument.

sem_timedwait() The timeout is based on the CLOCK_REALTIME clock.

send() The list of errors now includes EPIPE.

shm_open() Corrected the interpretation of the *name* argument.

sigaction() Corrected the example (it isn't safe to call *printf()* in a signal handler).

spawn(), *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*,
spawnv(), *spawnve()*, *spawnvp()*, *spawnvpe()*
You can now execute a shell script.

The child process's *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* are now calculated.

timer_create() Don't use SIGEV_INTR or SIGEV_UNBLOCK for the event type.

vsnprintf() Corrected the returned values.

What's new in QNX Neutrino 6.2

Significant changes:

- New content
- Deprecated content
- Errata

New Content

addrinfo	TCP/IP address information
<i>dirctl()</i>	Control an open directory
<i>freeaddrinfo()</i>	Free an address information structure
<i>freeifaddrs()</i>	Free an address information structure
<i>gai_strerror()</i>	Return the <i>getaddrinfo()</i> error code
<i>getaddrinfo()</i>	Get address information
<i>getdomainname()</i>	Get the domain name of the current host
<i>gethostbyname2()</i>	Get a network host entry, given a name
<i>getifaddrs()</i>	Get a network interface address
<i>hwi_find_item()</i>	Find an item in the hwi_item structure

<i>hwi_find_tag()</i>	Find a tag in the <i>hwi_item</i> structure
<i>hwi_off2tag()</i>	Return a pointer to the start of a tag in the <i>hwinfo</i> area of the system page
<i>hwi_tag2off()</i>	Return the offset from the start of the <i>hwinfo</i> area of the system page
ICMP6	Internet Control Message Protocol for IPv6
<i>if_freenameindex()</i>	Free dynamic memory allocated by <i>if_nameindex()</i>
<i>if_indextoname()</i>	Map an interface index to its name
<i>if_nameindex()</i>	Return a list of interfaces
<i>if_nametoindex()</i>	Map an interface name to its index
<i>ifaddrs()</i>	Structure that describes an Internet host
INET6	Internet Protocol version 6 family
<i>inet_net_ntop()</i>	Convert an Internet network number to CIDR format
<i>inet_net_pton()</i>	Convert an Internet network number from CIDR format to network format
IPv6	Internet Protocol version 6
IPsec	Internet security protocol
<i>mallinfo()</i>	Get memory allocation information
<i>mallopt()</i>	Control the memory allocation
<i>mcheck()</i>	Enable memory allocation routine consistency checks
<i>memalign()</i>	Allocate aligned memory
<i>mprobe()</i>	Perform consistency check on memory

<i>posix_memalign()</i>	Allocate aligned memory
<i>procmgr_session()</i>	Provide process manager session support
<i>_resmgr_handle_grow()</i>	Expand the capacity of the device manager database
<i>_resmgr_io_func()</i>	Retrieve an I/O function from an I/O function table
<i>resmgr_iofuncs()</i>	Extract the I/O function pointers associated with c connection
<i>_resmgr_ocb()</i>	Retrieve an Open Control Block
<i>sched_get_priority_adjust()</i>	Calculate the allowable priority for the scheduling policy
<i>seekdir()</i>	Set the position for the next read of the directory stream
<i>_sleepon_broadcast()</i>	Wake up multiple threads
<i>_sleepon_destroy()</i>	Destroy a sleepon
<i>_sleepon_init()</i>	Initialize a sleepon
<i>_sleepon_lock()</i>	Lock a sleepon
<i>_sleepon_signal()</i>	Wake up a single thread
<i>_sleepon_unlock()</i>	Unlock a sleepon

<i>_sleepon_wait()</i>	Wait on a sleep
<i>tcsetid()</i>	Make a terminal device a controlling device
<i>strtoimax()</i> , <i>strtoumax()</i>	Convert a string to an integer type
<i>telldir()</i>	Get the location associated with the directory stream
<i>valloc()</i>	Allocate a heap block aligned on a page boundary
<i>wcstoimax()</i> , <i>wcstoumax()</i>	Convert a wide-character string to an integer type

Deprecated Content

- *getpriority()* — use *getprio()* or *SchedGet()* instead.
- *setpriority()* — use *setprio()* or *SchedSet()* instead.

Errata

snprintf() Corrected the Returns section and Classifications

What's new in the QNX Neutrino 6.1.0 docs

Significant changes:

- New content
- Deprecated content

New content

The following functions have been added:

Wide-character functions

Wide-character versions of many functions

<i>InterruptHookTrace()</i>	Attach the pseudo interrupt handler that's used by the instrumented module
<i>iofdinfo()</i>	Retrieve server attributes
<i>iofunc_finfo()</i>	Handle an <code>_IO_FDINFO</code> message
<i>iofunc_finfo_default()</i>	Default handler for <code>_IO_FDINFO</code> messages
<i>MsgVerifyEvent()</i> , <i>MsgVerifyEvent_r()</i>	Check the validity of a receive ID and an event configuration
<i>resmgr_unbind()</i>	Remove an OCB
<i>straddstr()</i>	Concatenate one string on to the end of another
<i>SyncCtl()</i> , <i>SyncCtl_r()</i>	Perform an operation on a synchronization object
<i>SyncMutexEvent()</i> , <i>SyncMutexEvent_r()</i>	Attach an event to a mutex
<i>SyncMutexRevive()</i> , <i>SyncMutexRevive_r()</i>	Revive a mutex
<i>thread_pool_control()</i>	Control the thread pool behavior
<i>thread_pool_limits()</i>	Wrapper function for <i>thread_pool_control()</i>
<i>TraceEvent()</i>	Trace kernel events

Deprecated content

The following function has been deprecated:

matherr() Handle errors in math library functions



Summary of Functions



Summary of function categories

We've organized the functions in the C library into the following categories:

Asynchronous I/O functions

Asynchronous read, write, and other I/O operations.

Atomic functions

Thread-safe integer manipulation functions.

Character manipulation functions

Single-character functions for upper/lowercase conversions.

Conversion functions

Convert values from one representation to another (e.g. numeric values to strings).

Directory functions

Directory services (change, open, close, etc.).

Dispatch interface functions

Handle different event types, including messages, pulse codes, and signals.

File manipulation functions

File operations (change permissions, delete, rename, etc.)

IPC functions

Traditional InterProcess Communication functions.

Hardware functions

These functions work with PCI and other devices.

Math functions

Perform computations such as the common trigonometric calculations. These functions operate with floating-point values.

Memory allocation functions

Allocate and deallocate memory.

Memory manipulation functions

Manipulate blocks of memory.

Message queue functions

Nonblocking message-passing facilities.

Multibyte character functions

ANSI C functions for processing multibyte and wide characters.

QNX Neutrino-specific IPC functions

Native message-passing and related functions.

Operating system I/O functions

POSIX functions for performing I/O at a lower level than the C Language stream I/O functions (e.g. *fopen()*, *fread()*, *fwrite()*, and *fclose()*).

PC Card functions

Native PC Card functions.

Platform-specific functions

Invoke Intel 80x86 and other processor-related functions directly from a program.

Process environment functions

For process identification, user identification, process groups, system identification, system time and process time, environment variables, terminal identification, and configurable system variables.

Process manipulation functions

For process creation, execution, and termination; signal handling; and timer operations.

Realtime timer functions

Rich set of “inexpensive” timer functions that are quick to create and manipulate.

Resource manager functions

These functions help you create resource managers.

Searching and sorting functions

Perform various search and sort operations (do a binary search on a sorted array, find one string inside another, etc.).

Shared-memory functions

Create and manipulate shared-memory regions.

Signal functions Rich set of functions for handling and sending signals.

Stream I/O functions

The “standard” functions to read and write files. Data can be transmitted under format control or as characters, strings, or blocks of memory.

String manipulation functions

Manipulate a character string, i.e. an array of zero or more adjacent characters followed by a NUL character (`\0`) that marks the end of the string.

System database functions

Allow an application to access group and user database information.

System message log functions

This set of functions controls the system log.

TCP/IP functions

Handle TCP/IP network communications and the TCP/IP database files.

Terminal control functions

Set and control terminal attributes (baud rate, flow control, etc.).

Thread functions

Operate on threads and the objects used to synchronize threads.

Time functions Obtain and manipulate times and dates.

Variable-length argument list functions

Process a variable number of arguments to a function.

Wide-character functions

Wide-character versions of functions from other function summary categories.

The following subsections describe these function categories in more detail. Each function is noted with a brief description of its purpose.

Asynchronous I/O functions

These functions perform *asynchronous* read, write, and other I/O operations.



Asynchronous I/O operations aren't currently supported.

The following functions are defined:

<i> aio_cancel()</i>	Cancel an asynchronous I/O operation
<i> aio_error()</i>	Get the error status for an asynchronous I/O operation
<i> aio_fsync()</i>	Asynchronously synchronize a file
<i> aio_read()</i>	Asynchronously read from a file

<i>aio_return()</i>	Get the return status for an asynchronous I/O operation
<i>aio_suspend()</i>	Wait for asynchronous I/O operations to complete
<i>aio_write()</i>	Asynchronously write to a file

Atomic functions

These functions manipulate an integer in a thread-safe way. On a multiprocessor system, even a simple:

```
/*
   Assuming x is an unsigned variable shared between two
   or more threads or a thread and an interrupt handler.
*/
x ^= 0xdeadbeef;
```

may cause *x* to be in an undefined state if multiple threads running simultaneously on multiple processors execute this code at the same time.

Use the *atomic*()* functions to ensure that your integer operations are carried out properly:

```
atomic_toggle( &x, 0xdeadbeef );
```

<i>atomic_add()</i>	Safely add to a variable
<i>atomic_add_value()</i>	Safely add to a variable, returning the previous value
<i>atomic_clr()</i>	Safely clear a variable
<i>atomic_clr_value()</i>	Safely clear a variable, returning the previous value
<i>atomic_set()</i>	Safely set bits in a variable

<i>atomic_set_value()</i>	Safely set bits in a variable, returning the previous value
<i>atomic_sub()</i>	Safely subtract from a variable
<i>atomic_sub_value()</i>	Safely subtract from a variable, returning the previous value
<i>atomic_toggle()</i>	Safely toggle a variable
<i>atomic_toggle_value()</i>	Safely toggle a variable, returning the previous value

Character manipulation functions

These functions operate on single characters of type **char**. The functions test characters in various ways and convert them between upper and lowercase. (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

<i>isalnum()</i>	Test a character to see if it’s alphanumeric
<i>isalpha()</i>	Test to see if a character is a letter
<i>isascii()</i>	Test for a character in the range 0 to 127
<i>isctrl()</i>	Test a character to see if it’s a control character
<i>isdigit()</i>	Test for any decimal digit
<i>isgraph()</i>	Test for any printable character except space
<i>islower()</i>	Test for any lowercase letter
<i>isprint()</i>	Test for any printable character, including space

<i>ispunct()</i>	Test for any punctuation character
<i>isspace()</i>	Test for a whitespace character
<i>isupper()</i>	Test for any uppercase letter
<i>isxdigit()</i>	Test for any hexadecimal digit
<i>tolower()</i>	Convert a character to lowercase
<i>toupper()</i>	Convert a character to uppercase

Conversion functions

These functions perform conversions between objects of various types and strings:

<i>atof()</i>	Convert a string into a double
<i>atoh()</i>	Convert a string containing a hexadecimal number into an unsigned number
<i>atoi()</i>	Convert a string into an integer
<i>atol()</i> , <i>atoll()</i>	Convert a string into a long integer
<i>ENDIAN_BE16()</i>	Return a big-endian 16-bit value in native format
<i>ENDIAN_BE32()</i>	Return a big-endian 32-bit value in native format
<i>ENDIAN_BE64()</i>	Return a big-endian 64-bit value in native format
<i>ENDIAN_LE16()</i>	Return a little-endian 16-bit value in native format
<i>ENDIAN_LE32()</i>	Return a little-endian 32-bit value in native format

<i>ENDIAN_LE64()</i>	Return a little-endian 64-bit value in native format
<i>ENDIAN_RET16()</i>	Return an endian-swapped 16-bit value
<i>ENDIAN_RET32()</i>	Return an endian-swapped 32-bit value
<i>ENDIAN_RET64()</i>	Return an endian-swapped 64-bit value
<i>ENDIAN_SWAP16()</i>	Endian-swap a 16-bit value in place
<i>ENDIAN_SWAP32()</i>	Endian-swap a 32-bit value in place
<i>ENDIAN_SWAP64()</i>	Endian-swap a 64-bit value in place
<i>htonl()</i>	Convert a 32-bit value from host-byte order to network-byte order
<i>htons()</i>	Convert a 16-bit value from host-byte order to network-byte order
<i>itoa()</i>	Convert an integer into a string, using a given base
<i>ltoa()</i> , <i>lltoa()</i>	Convert a long integer value into a string, using a given base
<i>nsec2timespec()</i>	Convert nanoseconds to a timespec structure
<i>ntohl()</i>	Convert network-byte order value
<i>ntohs()</i>	Convert network-byte order value
<i>strtod()</i>	Convert a string into a double

<i>strtoimax()</i> , <i>strtoumax()</i>	Convert a string into an integer
<i>strtol()</i> , <i>strtoll()</i>	Convert a string into a long integer
<i>strtoul()</i> , <i>strtoull()</i>	Convert a string into an unsigned long integer
timespec	Time-specification structure
<i>timespec2nsec()</i>	Convert a timespec structure to nanoseconds
<i>ultoa()</i> , <i>lltoa()</i>	Convert an unsigned long integer into a string, using a given base
<i>UNALIGNED_PUT16()</i>	Write a misaligned 16-bit value safely
<i>UNALIGNED_PUT32()</i>	Write a misaligned 32-bit value safely
<i>UNALIGNED_PUT64()</i>	Write a misaligned 64-bit value safely
<i>UNALIGNED_RET16()</i>	Access a misaligned 16-bit value safely
<i>UNALIGNED_RET32()</i>	Access a misaligned 32-bit value safely
<i>UNALIGNED_RET64()</i>	Access a misaligned 64-bit value safely
<i>utoa()</i>	Convert an unsigned integer into a string, using a given base
<i>wordexp()</i>	Perform word expansions

wordfree() Free a word expansion buffer

See also the following functions, which convert the cases of characters and strings:

- *strlwr()*
- *strupr()*
- *tolower()*
- *toupper()*

Directory functions

These functions pertain to directory manipulation:

alphasort() Compare two directory entries

chdir() Change the current working directory

chroot() Change the root directory

closedir() Close a directory

dircntl() Control an open directory

dirname() Report the parent directory name of a file pathname

getcwd() Get the name of the current working directory

getwd() Get current working directory pathname

glob() Find paths matching a pattern

globfree() Free storage allocated by a call to *glob()*

mkdir() Create a subdirectory

mount() Mount a filesystem

mount_parse_generic_args()
Strip off common mount arguments

<i>opendir()</i>	Open a directory file
<i>pathfind()</i> , <i>pathfind_r()</i>	Search for a file in a list of directories
<i>readdir()</i>	Get information about the next matching filename
<i>readdir_r()</i>	Get information about the next matching filename
<i>realpath()</i>	Resolve a pathname
<i>rewinddir()</i>	Reset the position of a directory stream to the start of the directory
<i>rmdir()</i>	Delete an empty directory
<i>scandir()</i>	Scan a directory
<i>seekdir()</i>	Set the position for the next read of the directory stream
<i>telldir()</i>	Get the location associated with the directory stream
<i>umount()</i>	Unmount a filesystem

Dispatch interface functions

These functions make up the dispatch interface where you can handle different event types including messages, pulse codes, and signals. The functions cover dispatch contexts, attaching events, attaching pathnames and file descriptors to dispatch contexts, thread pools, etc. For an overview of these functions, see “Components of a resource manager” in the Writing a Resource Manager chapter of the QNX Neutrino *Programmer’s Guide*.

<i>dispatch_block()</i>	Block while waiting for an event
<i>dispatch_context_alloc()</i>	Return a dispatch context

<i>dispatch_context_free()</i>	Free a dispatch context
<i>dispatch_create()</i>	Allocate a dispatch handle
<i>dispatch_destroy()</i>	Destroy a dispatch handle
<i>dispatch_handler()</i>	Handle events received by <i>dispatch_block()</i>
<i>dispatch_timeout()</i>	Set a timeout
<i>dispatch_unblock()</i>	Unblock all of the threads that are blocked on a dispatch handle
<i>message_attach()</i>	Attach a message range
<i>message_connect()</i>	Create a connection to a channel
<i>message_detach()</i>	Detach a message range
<i>name_attach()</i>	Register a name in the namespace and create a channel
<i>name_detach()</i>	Remove a name from the namespace and destroy the channel
_pulse	Structure that describes a pulse
<i>pulse_attach()</i>	Attach a handler function to a pulse code
<i>pulse_detach()</i>	Detach a handler function from a pulse code

<i>resmgr_attach()</i>	Attach a path to a pathname space
<i>resmgr_block()</i>	Block while waiting for a message
resmgr_connect_funcs_t	Table of POSIX-level connect functions
<i>resmgr_context_alloc()</i>	Allocate a resource-manager context
<i>resmgr_context_free()</i>	Free a resource-manager context
resmgr_context_t	Context information that's passed between resource-manager functions
<i>resmgr_detach()</i>	Remove a pathname from the pathname space
<i>resmgr_devino()</i>	Get the device and inode number
<i>_resmgr_handle_grow()</i>	Expand the capacity of the device manager database
<i>resmgr_handler()</i>	Handle resource manager messages
<i>_resmgr_io_func()</i>	Retrieve an I/O function from an I/O function table
resmgr_io_funcs_t	Table of POSIX-level I/O functions
<i>resmgr_iofuncs()</i>	Extract the I/O function pointers associated with client connections
<i>resmgr_msgread()</i>	Read a message from a client

<i>resmgr_msgreadv()</i>	Read a message from a client
<i>resmgr_msgwrite()</i>	Write a message to a client
<i>resmgr_msgwritev()</i>	Write a message to a client
<i>_RESMGR_NPARTS()</i>	Get a given number of parts from the <i>ctp->iov</i> structure
<i>_resmgr_ocb()</i>	Retrieve an Open Control Block
<i>resmgr_open_bind()</i>	Associate an OCB with an open request
<i>resmgr_pathname()</i>	Return the pathname associated with an ID
<i>_RESMGR_PTR()</i>	Get one part from the <i>ctp->iov</i> structure and fill in its fields
<i>_RESMGR_STATUS()</i>	Set the status member of a resource-manager context
<i>resmgr_unbind()</i>	Remove an OCB
<i>select_attach()</i>	Attach a file descriptor to a dispatch handle
<i>select_detach()</i>	Detach a file descriptor from a dispatch handle
<i>select_query()</i>	Decode the last select event
<i>thread_pool_create()</i>	Create a thread pool handle

thread_pool_control()

Control the thread pool behavior

thread_pool_destroy()

Free the memory allocated to a thread pool

thread_pool_limits()

Wrapper function for *thread_pool_control()*

thread_pool_start()

Start a thread pool

File manipulation functions

These functions operate directly with files. The following functions are defined:

access() Check to see if a file or directory can be accessed

chmod() Change the permissions for a file

chown() Change the user ID and group ID of a file

eaccess() Check to see if a file or directory can be accessed (extended version)

glob() Find paths matching a pattern

globfree() Free storage allocated by a call to *glob()*

fchmod() Change the permissions for a file

fchown() Change the user ID and group ID of a file

fpathconf() Return the value of a configurable limit associated with a file or directory

ftruncate(), *ftruncate64()*

Truncate a file

futime() Record the modification time for a file

<i>lchown()</i>	Change the user ID and group ID of a file or symbolic link
<i>lstat()</i> , <i>lstat64()</i>	Get information about a file or directory
<i>ltruncate()</i>	Truncate a file at a given position
<i>mkfifo()</i>	Create a FIFO special file
<i>mkstemp()</i>	Make a unique temporary filename, and open the file
<i>mktemp()</i>	Make a unique temporary filename
<i>pathconf()</i>	Return the value of a configurable limit
<i>pclose()</i>	Close a pipe
<i>pwrite()</i> , <i>pwrite64()</i>	Write into a file without changing the file pointer
<i>remove()</i>	Remove a link to a file
<i>rename()</i>	Rename a file
<i>stat()</i> , <i>stat64()</i>	Get information about a file or directory, given a path
<i>statvfs()</i> , <i>statvfs64()</i>	Get filesystem information, given a path
<i>sync()</i>	Synchronize filesystem updates
<i>tempnam()</i>	Create a name for a temporary file
<i>truncate()</i>	Truncate a file to a specified length
<i>tmpnam()</i>	Generate a unique string for use as a filename
<i>unlink()</i>	Remove a link to a file
<i>utime()</i>	Record the modification time for a file or directory
<i>utimes()</i>	Set a file's access and modification times

IPC functions

These functions deal with InterProcess Communications.

<i>flock()</i>	Apply or remove an advisory lock on an open file
<i>lockf()</i>	Record locking on files
<i>mlock()</i>	Lock a buffer in physical memory
<i>mlockall()</i>	Lock a process's address space
<i>mmap()</i> , <i>mmap64()</i>	Map a memory region into a process address space
<i>mprotect()</i>	Change memory protection
<i>msync()</i>	Synchronize memory with physical storage
<i>munlock()</i>	Unlock a buffer
<i>munlockall()</i>	Unlock a process's address space
<i>munmap()</i>	Unmap previously mapped addresses
<i>pthread_barrier_destroy()</i>	Destroy a barrier object
<i>pthread_barrier_init()</i>	Initialize a barrier object
<i>pthread_barrier_wait()</i>	Synchronize at a barrier
<i>pthread_barrierattr_destroy()</i>	Destroy a barrier attributes object
<i>pthread_barrierattr_getpshared()</i>	Get the process-shared attribute of a barrier attributes object

pthread_barrierattr_init()

Initialize a barrier attributes object

pthread_barrierattr_setpshared()

Set the process-shared attribute of a barrier attributes object

pthread_cond_broadcast()

Unblock threads waiting on a condition

pthread_cond_destroy()

Destroy a condition variable

pthread_cond_init()

Initialize a condition variable

pthread_cond_signal()

Unblock the thread waiting on a condition variable

pthread_cond_timedwait()

Wait on a condition variable, with a time limit

pthread_cond_wait()

Wait on a condition variable

pthread_condattr_destroy()

Destroy a condition variable attribute object

pthread_condattr_getclock()

Get the clock attribute from a condition-variable attribute object

pthread_condattr_getpshared()

Get the process-shared attribute from a condition variable attribute object

pthread_condattr_init()

Initialize a condition variable attribute object

pthread_condattr_setclock()

Set the clock attribute in a condition-variable attribute object

pthread_condattr_setshared()

Set the process-shared attribute in a condition-variable attribute object

pthread_mutex_destroy()

Destroy a mutex

pthread_mutex_getprioceiling()

Get a mutex's priority ceiling

pthread_mutex_init()

Initialize a mutex

pthread_mutex_lock()

Lock a mutex

pthread_mutex_setprioceiling()

Set a mutex's priority ceiling

pthread_mutex_timedlock()

Lock a mutex

pthread_mutex_trylock()

Attempt to lock a mutex

pthread_mutex_unlock()

Unlock a mutex

pthread_mutexattr_destroy()

Destroy a mutex attribute object

pthread_mutexattr_getprioceiling()

Get the priority ceiling of a mutex attribute object

- pthread_mutexattr_getprotocol()*
Get a mutex's scheduling protocol
- pthread_mutexattr_getpshared()*
Get the process-shared attribute from a mutex attribute object
- pthread_mutexattr_getrecursive()*
Get the recursive attribute from a mutex attribute object
- pthread_mutexattr_gettype()*
Get a mutex type
- pthread_mutexattr_init()*
Initialize the mutex attribute object
- pthread_mutexattr_setprioceiling()*
Set the priority ceiling of a mutex attribute object
- pthread_mutexattr_setprotocol()*
Set a mutex's scheduling protocol
- pthread_mutexattr_setpshared()*
Set the process-shared attribute in mutex attribute object
- pthread_mutexattr_setrecursive()*
Set the recursive attribute in mutex attribute object
- pthread_mutexattr_settype()*
Set a mutex type
- pthread_once()* Dynamic package initialization
- pthread_rwlock_destroy()*
Destroy a read/write lock

pthread_rwlock_init()

Initialize a read/write lock

pthread_rwlock_rdlock()

Acquire a shared read lock on a read/write lock

pthread_rwlock_timedrdlock()

Lock a read-write lock for writing

pthread_rwlock_timedwrlock()

Attempt to acquire an exclusive write lock on a read/write lock

pthread_rwlock_tryrdlock()

Attempt to acquire a shared read lock on a read/write lock

pthread_rwlock_trywrlock()

Attempt to acquire an exclusive write lock on a read/write lock

pthread_rwlock_unlock()

Unlock a read/write lock

pthread_rwlock_wrlock()

Acquire an exclusive write lock on a read/write lock

pthread_rwlockattr_destroy()

Destroy a read-write lock attribute object

pthread_rwlockattr_getshared()

Get the process-shared attribute of a read-write lock attribute object

pthread_rwlockattr_init()

Create a read-write lock attribute object

<i>pthread_rwlockattr_setpshared()</i>	Set the process-shared attribute of a read-write lock attribute object
<i>pthread_spin_destroy()</i>	Destroy a thread spinlock
<i>pthread_spin_init()</i>	Initialize a thread spinlock
<i>pthread_spin_lock()</i>	Lock a thread spinlock
<i>pthread_spin_trylock()</i>	Try locking a thread spinlock
<i>pthread_spin_unlock()</i>	Unlock a thread spinlock
<i>readcond()</i>	Read data from a terminal device
<i>sem_close()</i>	Close a named semaphore
<i>sem_destroy()</i>	Destroy a semaphore
<i>sem_getvalue()</i>	Get the value of a semaphore (named or unnamed)
<i>sem_init()</i>	Initialize a semaphore
<i>sem_open()</i>	Create or access a named semaphore
<i>sem_post()</i>	Increment a semaphore
<i>sem_timedwait()</i>	Wait on a semaphore, with a timeout
<i>sem_trywait()</i>	Wait on a semaphore, but don't block
<i>sem_unlink()</i>	Destroy a named semaphore
<i>sem_wait()</i>	Wait on a semaphore
<i>sync()</i>	Synchronize filesystem updates

Hardware functions

These functions work with PCI and other devices for operations such as determining whether or not a PCI BIOS is present, attaching a driver to a PCI device, and so on.

The following functions are defined:

<i>pci_attach()</i>	Connect to the PCI server
<i>pci_attach_device()</i>	Attach a driver to a PCI device
<i>pci_detach()</i>	Disconnect from the PCI server
<i>pci_detach_device()</i>	Detach a driver from a PCI device
<i>pci_find_class()</i>	Find devices that have a specific Class Code
<i>pci_find_device()</i>	Find the PCI device with a given device ID and vendor ID
<i>pci_irq_routing_options()</i>	Retrieve PCI IRQ routing information
<i>pci_map_irq()</i>	Map an interrupt pin to an IRQ
<i>pci_present()</i>	Determine whether or not PCI BIOS is present
<i>pci_read_config()</i>	Read from the configuration space of a PCI device
<i>pci_read_config8()</i>	Read a byte from the configuration space of a device
<i>pci_read_config16()</i>	Read 16-bit values from the configuration space of a device

<i>pci_read_config32()</i>	Read 32-bit values from the configuration space of a device
<i>pci_rescan_bus()</i>	Rescan the PCI bus for added or removed devices
<i>pci_write_config()</i>	Write to the configuration space of a PCI device
<i>pci_write_config8()</i>	Write bytes to the configuration space of a PCI device
<i>pci_write_config16()</i>	Write 16-bit values to the configuration space of a device
<i>pci_write_config32()</i>	Write 32-bit values to the configuration space of a device
<i>hwi_find_item()</i>	Find an item in the <code>hwi_item</code> structure
<i>hwi_find_tag()</i>	Find a tag in the <code>hwi_item</code> structure
<i>hwi_off2tag()</i>	Return a pointer to the start of a tag in the <code>hwinfo</code> area of the system page
<i>hwi_tag2off()</i>	Return the offset from the start of the <code>hwinfo</code> area of the system page

Math functions

The math functions are arranged in the following categories:

- Absolute values
- Bessel functions
- Divisions, remainders, and modular arithmetic

- Floating-point settings
- Gamma functions
- Logarithms and exponentials
- Miscellaneous
- Pseudo-random numbers
- Roots and powers
- Rounding
- Trigonometric and hyperbolic functions

Absolute values

<i>abs()</i>	Return the absolute value of an integer
<i>cabs()</i> , <i>cabsf()</i>	Compute the absolute value of a complex number
<i>fabs()</i> , <i>fabsf()</i>	Compute the absolute value of a double number
<i>labs()</i>	Calculate the absolute value of a long integer

Bessel functions

<i>j0()</i> , <i>j0f()</i>	Compute a Bessel function of the first kind
<i>j1()</i> , <i>j1f()</i>	Compute a Bessel function of the first kind
<i>jn()</i> , <i>jnf()</i>	Compute a Bessel function of the first kind
<i>y0()</i> , <i>y0f()</i>	Compute a Bessel function of the second kind
<i>y1()</i> , <i>y1f()</i>	Compute a Bessel function of the second kind
<i>yn()</i> , <i>ynf()</i>	Compute a Bessel function of the second kind

Division, remainders, and modular arithmetic

- div()* Calculate the quotient and remainder of a division operation
- drem()*, *dremf()*
Compute the remainder of two numbers
- fmod()*, *fmodf()*
Compute a residue, using floating-point modular arithmetic
- ldiv()* Perform division on long integers
- modf()*, *modff()*
Break a number into integral and fractional parts
- remainder()*, *remainderf()*
Compute the floating point remainder

Floating-point settings

These functions set or get attributes of floating-point operations:

- fp_exception_mask()*
Get or set the current exception mask
- fp_exception_value()*
Get the value of the current exception registers
- fp_precision()* Set or get the current precision
- fp_rounding()* Set or get the current rounding

Gamma functions

- gamma()*, *gamma_r()*, *gammaf()*, *gammaf_r()*
Log gamma function
- lgamma()*, *lgamma_r()*, *lgammaf()*, *lgammaf_r()*
Log gamma function

Logarithms and exponentials

The following routines calculate logarithms and exponentials:

<i>exp()</i> , <i>expf()</i>	Compute the exponential function of a number
<i>expm1()</i> , <i>expm1f()</i>	Compute the exponential of a number, then subtract 1
<i>frexp()</i> , <i>frexpf()</i>	Break a floating-point number into a normalized fraction and an integral power of 2
<i>ilogb()</i> , <i>ilogbf()</i>	Compute the integral part of a logarithm
<i>ldexp()</i> , <i>ldexpf()</i>	Multiply a floating-point number by an integral power of 2
<i>log()</i> , <i>logf()</i>	Compute the natural logarithm of a number
<i>log10()</i> , <i>log10f()</i>	Compute the logarithm (base 10) of a number
<i>log1p()</i> , <i>log1pf()</i>	Compute $\log(1+x)$
<i>logb()</i> , <i>logbf()</i>	Compute the radix-independent exponent
<i>scalb()</i> , <i>scalbf()</i>	Load the exponent of a radix-independent floating-point number
<i>scalbn()</i> , <i>scalbnf()</i>	Compute the exponent of a radix-independent floating-point number
<i>significand()</i> , <i>significandf()</i>	Compute the “significant bits” of a floating-point number

Miscellaneous

<i>copysign()</i> , <i>copysignf()</i>	Copy the sign bit from one number to another
<i>erf()</i> , <i>erff()</i>	Compute the error function of a number
<i>erfc()</i> , <i>erfcf()</i>	Complementary error function
<i>finite()</i> , <i>finitef()</i>	Determine if a number is finite
<i>hypot()</i> , <i>hypotf()</i>	Calculate the length of the hypotenuse for a right-angled triangle
<i>isinf()</i> , <i>isinf()</i>	Test for infinity
<i>isnan()</i> , <i>isnanf()</i>	Test for not-a-number (NaN)
<i>max()</i>	Return the greater of two numbers
<i>min()</i>	Return the lesser of two numbers
<i>nextafter()</i> , <i>nextafterf()</i>	Compute the next representable double-precision floating-point number

Pseudo-random numbers

The math library includes several sets of functions that you can use to generate pseudo-random numbers.

The simplest family consists of:

<i>rand()</i>	Compute a sequence of pseudo-random integers
<i>rand_r()</i>	Compute a sequence of pseudo-random integers in a thread-safe manner
<i>srand()</i>	Start a new sequence of pseudo-random integers for <i>rand()</i>

This set of functions uses a nonlinear additive feedback random-number generator, using a state array:

<i>initstate()</i>	Initialize a pseudo-random number generator
<i>random()</i>	Generate a pseudo-random number from the default state
<i>setstate()</i>	Reset the state of a pseudo-random number generator
<i>srandom()</i>	Set the seed for a pseudo-random number generator

This set of functions uses 48-bit arithmetic to produce pseudo-random numbers of various types:

<i>drand48()</i>	Generate a pseudo-random double
<i>erand48()</i>	Generate a pseudo-random double in a thread-safe manner
<i>jrand48()</i>	Generate a pseudo-random signed long integer in a thread-safe manner
<i>lcong48()</i>	Initialize a sequence of pseudo-random numbers
<i>lrand48()</i>	Generate a pseudo-random nonnegative long integer
<i>mrand48()</i>	Generate a pseudo-random signed long integer
<i>nrand48()</i>	Generate a pseudo-random nonnegative long integer in a thread-safe manner
<i>seed48()</i>	Initialize a sequence of pseudo-random numbers
<i>srand48()</i>	Initialize a sequence of pseudo-random numbers

Roots and powers

<i>cbrt()</i> , <i>cbrtf()</i>	Compute the cube root of a number
<i>pow()</i> , <i>powf()</i>	Raise a number to a given power
<i>sqrt()</i> , <i>sqrtf()</i>	Calculate the nonnegative square root of a number

Rounding

<i>ceil()</i> , <i>ceilf()</i>	Round up a value to the next integer
<i>floor()</i> , <i>floorf()</i>	Round down a value to the next integer
<i>rint()</i> , <i>rintf()</i>	Round to the nearest integral value

Trigonometric and hyperbolic functions

<i>acos()</i> , <i>acosf()</i>	Compute the arccosine of an angle
<i>acosh()</i> , <i>acoshf()</i>	Compute the inverse hyperbolic cosine
<i>asin()</i> , <i>asinf()</i>	Compute the arcsine of an angle
<i>asinh()</i> , <i>asinhf()</i>	Compute the inverse hyperbolic sine
<i>atan()</i> , <i>atanf()</i>	Compute the arctangent of an angle
<i>atanh()</i> , <i>atanhf()</i>	Compute the inverse hyperbolic tangent
<i>atan2()</i> , <i>atan2f()</i>	Compute the arctangent, determining the quadrant
<i>cos()</i> , <i>cosf()</i>	Compute the cosine of an angle
<i>cosh()</i> , <i>coshf()</i>	Compute the hyperbolic cosine
<i>sin()</i> , <i>sinf()</i>	Calculate the sine of an angle
<i>sinh()</i> , <i>sinhf()</i>	Compute the hyperbolic sine
<i>tan()</i> , <i>tanf()</i>	Calculate the tangent of an angle
<i>tanh()</i> , <i>tanhf()</i>	Calculate the hyperbolic tangent

Memory allocation functions

These functions allocate and deallocate blocks of memory:

<i>alloca()</i>	Allocate automatic space from the stack
<i>_amblksiz</i>	The increment for the break pointer
<i>_btext</i>	The beginning of the text segment
<i>calloc()</i>	Allocate space for an array
<i>cfree()</i>	Free allocated memory
<i>_edata</i>	The end of the data segment, excluding BSS data
<i>_end</i>	The end of the data segment, including BSS data
<i>_etext</i>	The end of the text segment
<i>free()</i>	Deallocate a block of memory
<i>ftw()</i>	Walk a file tree
<i>longjmp()</i>	Restore the environment saved by <i>setjmp()</i>
<i>mallinfo()</i>	Get memory allocation information
<i>malloc()</i>	Allocate memory
<i>mallopt()</i>	Control the memory allocation
<i>mcheck()</i>	Enable memory allocation routine consistency checks
<i>memalign()</i>	Allocate aligned memory
<i>mprobe()</i>	Perform consistency check on memory
<i>posix_memalign()</i>	Allocate aligned memory
<i>realloc()</i>	Allocate, reallocate or free a block of memory

<i>sbrk()</i>	Set the allocation break value for a program
<i>_scalloc()</i>	Allocate space for an array
<i>setjmp()</i>	Save the calling environment, for use by <i>longjmp()</i>
<i>siglongjmp()</i>	Restore the signal mask for a process, if one was saved
<i>sigsetjmp()</i>	Save the environment, including the signal mask
<i>_sfree()</i>	Deallocate a block of memory
<i>_smalloc()</i>	Allocate memory in blocks
<i>_srealloc()</i>	Allocate, reallocate or free a block of memory
<i>valloc()</i>	Allocate a heap block aligned on a page boundary

Memory manipulation functions

These functions manipulate blocks of memory. In each case, the address of the memory block and its size is passed to the function. (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

<i>brk()</i>	Change the amount of space allocated for the calling process’s data segment
<i>bzero()</i>	Set the first part of an object to null bytes
<i>ffs()</i>	Find the first bit set in a bit string
<i>index()</i>	Find a character in a string
<i>memccpy()</i>	Copy bytes until a given character is found
<i>memchr()</i>	Find the first occurrence of a character in a buffer
<i>memcmp()</i>	Compare a given number of characters in two objects

<i>memcpy()</i>	Copy a number of characters from one buffer to another
<i>memcpyv()</i>	Copy a given number of structures
<i>memcmp()</i>	Compare a given number of characters of two objects, without case sensitivity
<i>mem_offset()</i> , <i>mem_offset64()</i>	Find offset of a mapped typed memory block
<i>memmove()</i>	Copy bytes from one buffer to another, handling overlapping memory correctly
<i>memset()</i>	Set the first part of an object to a given value
<i>mlock()</i>	Lock a buffer in physical memory
<i>mlockall()</i>	Lock a process's address space
<i>mmap()</i> , <i>mmap64()</i>	Map a memory region into a process address space
<i>mmap_device_io()</i>	Gain access to a device's registers
<i>mmap_device_memory()</i>	Map a device's physical memory into a process's address space
<i>msync()</i>	Synchronize memory with physical storage
<i>munlock()</i>	Unlock a buffer
<i>munlockall()</i>	Unlock a process's address space
<i>munmap_device_io()</i>	Free access to a device's registers
<i>munmap_device_memory()</i>	Unmap previously mapped addresses

<i>posix_mem_offset()</i> , <i>posix_mem_offset64()</i>	Find offset and length of a mapped typed memory block
<i>rindex()</i>	Find the last occurrence of a character in a string
<i>shm_ctl()</i>	Give special attributes to a shared memory object
<i>swab()</i>	Endian-swap a given number of bytes

See the section “String manipulation functions” for descriptions of functions that manipulate strings of data.

Message queue functions

These functions deal with message queues:

<i>mq_close()</i>	Close a message queue
<i>mq_getattr()</i>	Get a message queues attributes
<i>mq_notify()</i>	Ask to be notified when there’s a message in the queue
<i>mq_open()</i>	Open a message queue
<i>mq_receive()</i>	Receive a message from a queue
<i>mq_send()</i>	Send a message to a queue
<i>mq_setattr()</i>	Set a queue’s attributes
<i>mq_timedreceive()</i>	Receive a message from a message queue
<i>mq_timedsend()</i>	Send a message to a message queue
<i>mq_unlink()</i>	Remove a queue

Multibyte character functions

These ANSI C functions provide capabilities for processing multibyte characters. (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

<i>mblen()</i>	Count the bytes in a multibyte character
<i>mbrlen()</i>	Count the bytes in a multibyte character (restartable)
<i>mbrtowc()</i>	Convert a multibyte character into a wide character (restartable)
<i>mbstowcs()</i>	Determine the status of the conversion object used for restartable <i>mb*()</i> functions
<i>mbsrtowcs()</i>	Convert a multibyte-character string into a wide-character string (restartable)
<i>mbstowcs()</i>	Convert a multibyte-character string into a wide-character string
<i>mbtowc()</i>	Convert a multibyte character into a wide character

QNX Neutrino-specific IPC functions

The following functions are defined:

<i>ChannelCreate()</i> , <i>ChannelCreate_r()</i>	Create a communications channel
<i>ChannelDestroy()</i> , <i>ChannelDestroy_r()</i>	Destroy a communications channel
<i>_msg_info</i>	Information about a message
<i>MsgDeliverEvent()</i> , <i>MsgDeliverEvent_r()</i>	Deliver an event through a channel
<i>MsgError()</i> , <i>MsgError_r()</i>	Unblock a client and set its <i>errno</i>

MsgInfo(), *MsgInfo_r()*

Get additional information about a message

MsgKeyData(), *MsgKeyData_r()*

Pass data through a common client

MsgRead(), *MsgRead_r()*

Read data from a message

MsgReadv(), *MsgReadv_r()*

Read data from a message

MsgReceive(), *MsgReceive_r()*

Wait for a message or pulse on a channel

MsgReceivePulse(), *MsgReceivePulse_r()*

Receive a pulse on a channel

MsgReceivePulsev(), *MsgReceivePulsev_r()*

Receive a pulse on a channel

MsgReceivev(), *MsgReceivev_r()*

Wait for a message or pulse on a channel

MsgReply(), *MsgReply_r()*

Reply with a message

MsgReplyv(), *MsgReplyv_r()*

Reply with a message

MsgSend(), *MsgSend_r()*

Send a message to a channel

MsgSendnc(), *MsgSendnc_r()*

Send a message to a channel (non-cancellation point)

MsgSendPulse(), *MsgSendPulse_r()*

Send a pulse to a channel

MsgSendsv(), *MsgSendsv_r()*

Send a message to a channel

MsgSendsvnc(), *MsgSendsvnc_r()*

Send a message to a channel (non-cancellation point)

MsgSendv(), *MsgSendv_r()*

Send a message to a channel

MsgSendvnc(), *MsgSendvnc_r()*

Send a message to a channel (non-cancellation point)

MsgSendvs(), *MsgSendvs_r()*

Send a message to a channel

MsgSendvsnc(), *MsgSendvsnc_r()*

Send a message to a channel (non-cancellation point)

MsgVerifyEvent(), *MsgVerifyEvent_r()*

Check the validity of a receive ID and an event configuration

MsgWrite(), *MsgWrite_r()*

Write a reply

MsgWritev(), *MsgWritev_r()*

Write a reply

name_close() Close the file descriptor returned by *name_open()*

name_open() Open a name for a server connection

sigevent Structure that describes an event

SyncTypeCreate(), *SyncTypeCreate_r()*

Create a synchronization object

Operating system I/O functions

These functions operate at the operating-system level, and are included for compatibility with other C implementations. For new programs, we recommended that you use the functions in the section “File manipulation functions”, functions are defined portably and are part of the ANSI standard for the C language.

The functions in this section reference opened files and devices using a *file descriptor* that’s returned when the file is opened. The file descriptor is passed to the other functions.

The following functions are defined:

<i>chsize()</i>	Change the size of a file
<i>cfgopen()</i>	Open a configuration file
<i>close()</i>	Close a file at the operating system level
<i>_cmdfd()</i>	Return a file descriptor for the executable file
<i>_cmdname()</i>	Find the path used to invoke the current process
<i>creat()</i> , <i>creat64()</i>	Create and open a file at the operating system level
<i>devctl()</i>	Control a device
<i>dup()</i>	Duplicate a file descriptor
<i>dup2()</i>	Duplicate a file descriptor, specifying the new descriptor
<i>eof()</i>	Determine if the end-of-file has been reached
<i>fcfgopen()</i>	Open a configuration file
<i>fcntl()</i>	Provide control over an open file
<i>fdatasync()</i>	Write queued file data to disk

<i>fileno()</i>	Return the number of the file descriptor for a stream
<i>flink()</i>	Assign a pathname to a file descriptor
<i>flockfile()</i>	Acquire ownership of a file
<i>fstat(), fstat64()</i>	Obtain information about an open file, given a file descriptor
<i>fstatvfs(), fstatvfs64()</i>	Get filesystem information, given a file descriptor
<i>fsync()</i>	Write queued file and filesystem data to disk
<i>ftrylockfile()</i>	Acquire ownership of a file, without blocking
<i>funlockfile()</i>	Release ownership of a file
<i>GETIOVBASE()</i>	Get the base member of an <i>iov_t</i> structure
<i>GETIOVLEN()</i>	Get the length member of an <i>iov_t</i> structure
<i>getdtablesize()</i>	Get the size of the file descriptor table
<i>getrusage()</i>	Get information about resource utilization
<i>in8()</i>	Read an 8-bit value from a port
<i>in8s()</i>	Read 8-bit values from a port
<i>in16(), inbe16(), inle16()</i>	Read a 16-bit value from a port
<i>in16s()</i>	Read 16-bit values from a port
<i>in32(), inbe32(), inle32()</i>	Read a 32-bit value from a port
<i>in32s()</i>	Read 32-bit values from a port
<i>ioctl()</i>	Control device

<i>link()</i>	Create a link to an existing file
<i>lseek()</i> , <i>lseek64()</i>	Set the current file position at the OS level
<i>lio_listio()</i>	Initiate a list of I/O requests
<i>mknod()</i>	Make a new filesystem entry point
<i>modem_open()</i>	Open a serial port
<i>modem_read()</i>	Read bytes from a file descriptor
<i>modem_script()</i>	Run a script on a device
<i>modem_write()</i>	Write a string to a device
<i>name_close()</i>	Close the file descriptor obtained with <i>name_open()</i>
<i>name_open()</i>	Open a name for a server connection
<i>open()</i> , <i>open64()</i>	Open a file
<i>openfd()</i>	Open for private access a file associated with a given descriptor
<i>out8()</i>	Write an 8-bit value to a port
<i>out8s()</i>	Write 8-bit values to a port
<i>out16()</i> , <i>outbe16()</i> , <i>outle16()</i>	Write a 16-bit value to a port
<i>out16s()</i>	Write 16-bit values to a port
<i>out32()</i> , <i>outbe32()</i> , <i>outle32()</i>	Write a 32-bit value to a port
<i>out32s()</i>	Write 32-bit values to a port

<i>pathmgr_symlink()</i>	Create a symlink in the process manager
<i>pathmgr_unlink()</i>	Remove the link created by <i>pathmgr_symlink()</i>
<i>pipe()</i>	Create a pipe
<i>popen()</i>	Execute a command, creating a pipe to it
<i>pread(), pread64()</i>	Read from a file without moving the file pointer
<i>rdchk()</i>	Check to see if a read is likely to succeed
<i>read()</i>	Read bytes from a file
<i>readblock()</i>	Read blocks of data from a file
<i>readcond()</i>	Read data from a terminal device
<i>readlink()</i>	Place the contents of a symbolic link into a buffer
<i>readv()</i>	Read bytes from a file
<i>select()</i>	Check for files that are ready for reading or writing
<i>SETIOV()</i>	Fill in the fields of an <i>iov_t</i> structure
<i>sopen()</i>	Open a file for shared access
<i>sopenfd()</i>	Open for shared access a file associated with a given descriptor
<i>symlink()</i>	Create a symbolic link to a path
<i>tcischars()</i>	Determine the number of characters waiting to be read
<i>tell(), tell64()</i>	Determine the current file position
<i>umask()</i>	Set the file mode creation mask for the process

<i>uname()</i>	Get information about the operating system
<i>unlink()</i>	Delete a file
<i>write()</i>	Write bytes to a file
<i>writeblock()</i>	Write blocks of data to a file
<i>writenv()</i>	Write bytes to a file

PC Card functions

The following functions are defined:

<i>pccard_arm()</i>	Arm the devp-pccard server
<i>pccard_attach()</i>	Attach to the devp-pccard server
<i>pccard_detach()</i>	Detach from the devp-pccard server
<i>pccard_info()</i>	Obtain socket information from the devp-pccard server
<i>pccard_lock()</i>	Lock the window of the card in the selected socket
<i>pccard_raw_read()</i>	Read the raw CIS data from the PC Card
<i>pccard_unlock()</i>	Unlock the window of the card in the selected socket

Platform-specific functions

These functions are for invoking Intel 80x86 and other processor-related functions directly from a program. Functions that apply to the Intel 8086 CPU apply to that family including the 80286, 80386, 80486 and Pentium processors.

You'll also find endian-related functions listed here.

The following functions are defined:

ENDIAN_BE16()

Return a big-endian 16-bit value in native format

ENDIAN_BE32()

Return a big-endian 32-bit value in native format

ENDIAN_BE64()

Return a big-endian 64-bit value in native format

ENDIAN_LE16()

Return a little-endian 16-bit value in native format

ENDIAN_LE32()

Return a little-endian 32-bit value in native format

ENDIAN_LE64()

Return a little-endian 64-bit value in native format

ENDIAN_RET16()

Return an endian-swapped 16-bit value

ENDIAN_RET32()

Return an endian-swapped 32-bit value

ENDIAN_RET64()

Return an endian-swapped 64-bit value

ENDIAN_SWAP16()

Endian-swap a 16-bit value in place

ENDIAN_SWAP32()

Endian-swap a 32-bit value in place

ENDIAN_SWAP64()

Endian-swap a 64-bit value in place

_intr_v86()

Execute a real-mode software interrupt

offsetof()

Return the offset of an element within a structure

sysmgr_reboot()

Reboot a QNX Neutrino system

Process environment functions

These functions deal with process identification, user identification, process groups, system identification, system time and process time, environment variables, terminal identification, and configurable system variables:

<i>_argc</i>	The number of arguments passed to <i>main()</i>
<i>_argv</i>	A pointer to the vector of arguments passed to <i>main()</i>
<i>_auxv</i>	A pointer to a vector of auxiliary arguments to <i>main()</i>
<i>clearenv()</i>	Clear the process environment area
<i>confstr()</i>	Get configuration-defined string values
<i>ctermid()</i>	Generate the pathname of the current controlling terminal
<i>endutent()</i>	Close the current user-information file
<i>environ</i>	Pointer to the process's environment variables
<i>err(), errx()</i>	Display a formatted error message, and then exit
<i>errno</i>	Global error variable
<i>getegid()</i>	Get the effective group ID
<i>getenv()</i>	Get the value of an environment variable
<i>geteuid()</i>	Get the effective user ID
<i>getgid()</i>	Get the group ID
<i>getgrouplist()</i>	Determine the group access list for a user

<i>getgroups()</i>	Get the supplementary group IDs of the calling process
<i>getlogin()</i>	Get the user name associated with the calling process
<i>getlogin_r()</i>	Get the user name associated with the calling process
<i>getopt()</i>	Parse options from a command line
<i>getpgid()</i>	Get a process's group ID
<i>getpgrp()</i>	Get the process group
<i>getpid()</i>	Get the process ID
<i>getppid()</i>	Get the parent process ID
<i>getsid()</i>	Get the session ID of a process
<i>getuid()</i>	Get the user ID
<i>getutent()</i>	Read the next entry from the user-information file
<i>getutid()</i>	Search for an entry in the user-information file
<i>getutline()</i>	Get an entry from the user-information file
<i>initgroups()</i>	Initialize the supplementary group access list
<i>isatty()</i>	Test to see if a file descriptor is associated with a terminal
<i>login_tty()</i>	Prepare for a login in a tty
<i>main()</i>	The function where program execution begins
<i>ND_NODE_CMP()</i>	Compare two node descriptor IDs
<i>netmgr_ndtostr()</i>	Convert a node descriptor into a string

<i>netmgr_remote_nd()</i>	Get a node descriptor that's relative to a remote node
<i>netmgr_strtond()</i>	Convert a string into a node descriptor
<i>__progname</i>	The basename of the program being executed
<i>putenv()</i>	Add, change, or delete an environment variable
<i>pututline()</i>	Write an entry in the user-information file
<i>searchenv()</i>	Search the directories listed in an environment variable
<i>setegid()</i>	Set the effective group ID for a process
<i>setenv()</i>	Set one or more environment variables
<i>seteuid()</i>	Set the effective user ID
<i>setgid()</i>	Set the real, effective and saved group IDs
<i>setgroups()</i>	Set supplementary group IDs
<i>setlocale()</i>	Set a program's locale.
<i>setpgid()</i>	Join or create a process group
<i>setpgrp()</i>	Set the process group
<i>setregid()</i>	Set real and effective group IDs
<i>setreuid()</i>	Set real and effect user IDs
<i>setsid()</i>	Create a new session
<i>setuid()</i>	Set the real, effective and saved user IDs
<i>setutent()</i>	Return to the beginning of the user-information file
<i>strerror()</i>	Convert an error number into an error message

<i>sysconf()</i>	Return the value of a configurable system limit
<i>ttyname()</i>	Get a fully qualified pathname for a file
<i>ttyname_r()</i>	Get a fully qualified pathname for a file
<i>unsetenv()</i>	Remove an environment variable
utmp	Entry in a user-information file
<i>utmpname()</i>	Change the name of the user-information file
<i>verr()</i> , <i>verrx()</i>	Display a formatted error message, and then exit (varargs)
<i>vwarn()</i> , <i>vwarnx()</i>	Formatted error message (varargs)
<i>warn()</i> , <i>warnx()</i>	Formatted error message

Process manipulation functions

These functions deal with: process creation, execution, and termination; signal handling; and timer operations.

When you start a new process, it replaces the existing process if:

- You specify `P_OVERLAY` when calling one of the *spawn** functions.
- You call one of the *exec** routines.

The existing process may be suspended while the new process executes (control continues at the point following the place where the new process was started) in the following situations:

- You specify `P_WAIT` when calling one of the *spawn** functions.
- You call *system()*.

The following functions are defined:

<i>abort()</i>	Raise the SIGABRT signal to terminate program execution
<i>alarm()</i>	Schedule an alarm
<i>assert()</i>	Print a diagnostic message and optionally terminate the program
<i>atexit()</i>	Register functions to be called when the program terminates normally
<i>ConnectAttach()</i> , <i>ConnectAttach_r()</i>	Establish a connection between a process and a channel
<i>ConnectClientInfo()</i> , <i>ConnectClientInfo_r()</i>	Store information about a client connection
<i>ConnectDetach()</i> , <i>ConnectDetach_r()</i>	Break a connection between a process and a channel
<i>ConnectFlags()</i> , <i>ConnectFlags_r()</i>	Modify the flags associated with a connection
<i>ConnectServerInfo()</i> , <i>ConnectServerInfo_r()</i>	Store information about a connection
<i>daemon()</i>	Run a program in the background
<i>DebugBreak()</i>	Enter the process debugger
<i>DebugKDBreak()</i>	Enter the kernel debugger
<i>DebugKDOutput()</i>	Print text with the kernel debugger
<i>delay()</i>	Suspend a process for a given length of time
<i>dladdr()</i>	Translate an address to symbolic information

<i>dlclose()</i>	Close a shared object
<i>dlderror()</i>	Get dynamic loading diagnostic information
<i>dlopen()</i>	Gain access to an executable object file
<i>dlsym()</i>	Get the address of a symbol in a shared object
<i>execl()</i>	Execute a file
<i>execle()</i>	Execute a file
<i>execlp()</i>	Execute a file
<i>execlpe()</i>	Execute a file
<i>execv()</i>	Execute a file
<i>execve()</i>	Execute a file
<i>execvp()</i>	Execute a file
<i>execvpe()</i>	Execute a file
<i>_exit()</i>	Terminate the program
<i>exit()</i>	Terminate the program
<i>fork()</i>	Create a new process
<i>forkpty()</i>	Create a new process operating in a pseudo-tty
<i>getrlimit(), getrlimit64()</i>	Get the limit on a system resource
<i>getprio()</i>	Get the priority of a given process
<i>InterruptAttach(), InterruptAttach_r()</i>	Attach an interrupt handler to an interrupt source
<i>InterruptAttachEvent(), InterruptAttachEvent_r()</i>	Attach an event to an interrupt source

<i>InterruptDetach()</i> , <i>InterruptDetach_r()</i>	Detach an interrupt handler by ID
<i>InterruptDisable()</i>	Disable hardware interrupts
<i>InterruptEnable()</i>	Enable hardware interrupts
<i>InterruptHookIdle()</i>	Attach an “idle” interrupt handler
<i>InterruptHookTrace()</i>	Attach the pseudo interrupt handler that the instrumented module uses
<i>InterruptLock()</i>	Protect critical sections of an interrupt handler
<i>InterruptMask()</i>	Disable a hardware interrupt
<i>InterruptUnlock()</i>	Release a critical section locked with <i>InterruptLock()</i>
<i>InterruptUnmask()</i>	Enable a hardware interrupt
<i>InterruptWait()</i> , <i>InterruptWait_r()</i>	Wait for a hardware interrupt
<i>_intr_v86()</i>	Execute a real-mode software interrupt
<i>kill()</i>	Send a signal to a process or a group of processes
<i>killpg()</i>	Send a signal to a process group
<i>nap()</i>	Sleep for a given number of milliseconds
<i>napms()</i>	Sleep for a given number of milliseconds
<i>nice()</i>	Change the priority of a process

<i>openpty()</i>	Find an available pseudo-tty
<i>pause()</i>	Suspend the process until delivery of a signal
<i>procmgr_daemon()</i>	Run a process in the background
<i>procmgr_event_notify()</i>	Ask to be notified of system-wide events
<i>procmgr_event_trigger()</i>	Trigger a global system event
<i>procmgr_guardian()</i>	Let a daemon process takeover as parent = guardian
<i>procmgr_session()</i>	Provide process manager session support
<i>raise()</i>	Signal an exceptional condition
<i>SchedGet()</i> , <i>SchedGet_r()</i>	Get the scheduling policy for a thread
<i>SchedInfo()</i> , <i>SchedInfo_r()</i>	Get scheduler information
<i>SchedSet()</i> , <i>SchedSet_r()</i>	Set the scheduling policy for a thread
<i>SchedYield()</i> , <i>SchedYield_r()</i>	Yield to other threads
<i>setitimer()</i>	Set the value of an interval timer
<i>setprio()</i>	Set the priority of a process
<i>setrlimit()</i> , <i>setrlimit64()</i>	Set the limit on a system resource

<i>sigaction()</i>	Examine or specify the action associated with a signal
<i>sigaddset()</i>	Add a signal to a set
<i>sigblock()</i>	Add to the mask of signals to block
<i>sigdelset()</i>	Delete a signal from a set
<i>sigemptyset()</i>	Initialize a set to contain no signals
<i>sigfillset()</i>	Initialize a set to contain all signals
<i>sigismember()</i>	See if a given signal is in a given set
<i>sigmask()</i>	Construct a mask for a signal number
<i>signal()</i>	Set handling for exceptional conditions
<i>SignalAction()</i> , <i>SignalAction_r()</i>	Examine and/or specify actions for signals
<i>SignalKill()</i> , <i>SignalKill_r()</i>	Send a signal to a process group, process or thread
<i>SignalProcmask()</i> , <i>SignalProcmask_r()</i>	Modify or examine the signal blocked mask of a thread
<i>SignalSuspend()</i> , <i>SignalSuspend_r()</i>	Suspend a process until a signal is received
<i>SignalWaitinfo()</i> , <i>SignalWaitinfo_r()</i>	Select a pending signal
<i>sigpause()</i>	Wait for a signal
<i>sigpending()</i>	Examine the set of pending, masked signals for a process
<i>sigprocmask()</i>	Examine or change the signal mask for a process

<i>sigqueue()</i>	Queue a signal to a process
<i>sigsetmask()</i>	Set the mask of signals to block
<i>sigsuspend()</i>	Replace the signal mask, and then suspend the process
<i>sigtimedwait()</i>	Wait for a signal or a timeout
<i>sigunblock()</i>	Unblock signals
<i>sigwait()</i>	Wait for a pending signal
<i>sigwaitinfo()</i>	Wait for a pending signal and get its information
<i>sleep()</i>	Suspend a process for a given length of time
<i>spawn()</i>	Create and execute a new child process
<i>spawnl()</i>	Create and execute a new child process
<i>spawnle()</i>	Create and execute a new child process
<i>spawnlp()</i>	Create and execute a new child process
<i>spawnlpe()</i>	Create and execute a new child process
<i>spawnp()</i>	Create and execute a new child process
<i>spawnv()</i>	Create and execute a new child process
<i>spawnve()</i>	Create and execute a new child process
<i>spawnvp()</i>	Create and execute a new child process
<i>spawnvpe()</i>	Create and execute a new child process
<i>SyncCondvarSignal()</i> , <i>SyncCondvarSignal_r()</i>	Wake up any threads that are blocked on a synchronization object
<i>SyncCondvarWait()</i> , <i>SyncCondvarWait_r()</i>	Block a thread on a synchronization object

<i>SyncCtl()</i> , <i>SyncCtl_r()</i>	Perform an operation on a synchronization object
<i>SyncDestroy()</i> , <i>SyncDestroy_r()</i>	Destroy a synchronization object
<i>SyncMutexEvent()</i> , <i>SyncMutexEvent_r()</i>	Attach an event to a mutex
<i>SyncMutexLock()</i> , <i>SyncMutexLock_r()</i>	Lock a mutex synchronization object
<i>SyncMutexUnlock()</i> , <i>SyncMutexUnlock_r()</i>	Unlock a mutex synchronization object
<i>SyncMutexRevive()</i> , <i>SyncMutexRevive_r()</i>	Revive a mutex that's in the DEAD state
<i>SyncSemPost()</i> , <i>SyncSemPost_r()</i>	Increment a semaphore
<i>SyncSemWait()</i> , <i>SyncSemWait_r()</i>	Wait on a semaphore
<i>system()</i>	Execute a system command
<i>SYSPAGE_CPU_ENTRY()</i>	Return a CPU-specific entry from the system page
<i>SYSPAGE_ENTRY()</i>	Return an entry from the system page
<i>_syspage_ptr</i>	A pointer to the system page
<i>ThreadCancel()</i> , <i>ThreadCancel_r()</i>	Cancel a thread
<i>ThreadCreate()</i> , <i>ThreadCreate_r()</i>	Create a new thread

<i>ThreadCtl()</i> , <i>ThreadCtl_r()</i>	Control a thread
<i>ThreadDestroy()</i> , <i>ThreadDestroy_r()</i>	Destroy a thread immediately
<i>ThreadDetach()</i> , <i>ThreadDetach_r()</i>	Detach a thread from a process
<i>ThreadJoin()</i> , <i>ThreadJoin_r()</i>	Block until a thread terminates
<i>TraceEvent()</i>	Trace kernel events
<i>ualarm()</i>	Schedule an alarm
<i>usleep()</i>	Suspend a thread for a given number of microseconds
<i>vfork()</i>	Spawn a new process and block the parent
<i>wait()</i>	Wait for the status of a terminated child process
<i>wait3()</i>	Wait for a child process to change state
<i>wait4()</i>	Wait for a child process to terminate or stop
<i>waitid()</i>	Wait for a child process to change state
<i>waitpid()</i>	Suspend the calling process

There are eight *spawn*()* and *exec*()* functions each. The * is one to three letters, where:

- **l** or **v** (one is required) indicates the way the process parameters are passed
- **p** (optional) indicates that the **PATH** environment variable is searched to locate the program for the process
- **e** (optional) indicates that the environment variables are being passed

Realtime timer functions

These functions provide realtime timer capabilities:

<i>clock_getres()</i>	Get the resolution of the clock
<i>clock_gettime()</i>	Get the current time of a clock
<i>clock_nanosleep()</i>	High resolution sleep with specifiable clock
<i>clock_settime()</i>	Set a clock
<i>getitimer()</i>	Get the value of an interval timer
<i>nanosleep()</i>	Suspend process until a timeout or signal occurs
<i>nanospin()</i>	Busy-wait without thread blocking for a period of time
<i>nanospin_calibrate()</i>	Calibrate before calling <i>nanospin*()</i>
<i>nanospin_count()</i>	Busy-wait without blocking for a number of iterations
<i>nanospin_ns()</i>	Busy-wait without blocking for a period of time
<i>nanospin_ns_to_count()</i>	Convert a time in nanoseconds into a number of iterations
<i>sched_getparam()</i>	Get the current priority of a process
<i>sched_get_priority_adjust()</i>	Calculate the allowable priority for the scheduling policy

<i>sched_get_priority_max()</i>	Get the maximum value for the scheduling policy
<i>sched_get_priority_min()</i>	Get the minimum value for the scheduling policy
<i>sched_getscheduler()</i>	Get the current scheduling policy for a process
sched_param	Structure that describes scheduling parameters
<i>sched_rr_get_interval()</i>	Get the execution time limit of a process
<i>sched_setparam()</i>	Change the priority of a process
<i>sched_setscheduler()</i>	Change the priority and scheduling policy of a process
<i>sched_yield()</i>	Yield to other READY processes at the same priority
<i>timer_create()</i>	Create a timer
<i>timer_delete()</i>	Delete a timer
<i>timer_getexpstatus()</i>	Get the expiry status of a timer
<i>timer_getoverrun()</i>	Return the number of timer overruns
<i>timer_gettime()</i>	Get the amount of time left on a timer
<i>timer_settime()</i>	Set the expiration time for a timer

Resource manager functions

These functions help you create resource managers. For an overview of these functions, see “Components of a resource manager” in the Writing a Resource Manager chapter of the QNX Neutrino *Programmer’s Guide*.

`_io_connect` Structure of a resource manager’s connect message

`_io_connect_fstype_reply`

Structure of a connect message giving a status and a file type

`_io_connect_link_reply`

Structure of a connect message that redirects a client to another resource

iofinfo() Retrieve server attributes

iofunc_attr_init()

Initialize the default attribute structure

iofunc_attr_lock()

Lock the attribute structure

`iofunc_attr_t`

Attribute structure

iofunc_attr_trylock()

Try to lock the attribute structure

iofunc_attr_unlock()

Unlock the attribute structure

iofunc_check_access()

Check access permissions

iofunc_chmod() Handle an `_IO_CHMOD` message

<i>iofunc_chmod_default()</i>	Default handler for <code>_IO_CHMOD</code> messages
<i>iofunc_chown()</i>	Handle an <code>_IO_CHOWN</code> message
<i>iofunc_chown_default()</i>	Default handler for <code>_IO_CHOWN</code> messages
<i>iofunc_client_info()</i>	Return information about a client connection
<i>iofunc_close_dup()</i>	Frees all locks allocated for the client process
<i>iofunc_close_dup_default()</i>	Default handler for <code>_IO_CLOSE</code> messages
<i>iofunc_close_ocb()</i>	Return the memory allocated for an OCB
<i>iofunc_close_ocb_default()</i>	Return the memory allocated for an OCB
<i>iofunc_devctl()</i>	Handle an <code>_IO_DEVCTL</code> message
<i>iofunc_devctl_default()</i>	Default handler for <code>_IO_DEVCTL</code> messages
<i>iofunc_fdinfo()</i>	Handle an <code>_IO_FDINFO</code> message
<i>iofunc_fdinfo_default()</i>	Default handler for <code>_IO_FDINFO</code> messages
<i>iofunc_func_init()</i>	Initialize the default POSIX-layer function tables
<i>iofunc_link()</i>	Link two directories
<i>iofunc_lock()</i>	Lock a resource

<i>iofunc_lock_alloc()</i>	Allocate memory to lock structures
<i>iofunc_lock_default()</i>	Default handler for <code>_IO_LOCK</code> messages
<i>iofunc_lock_free()</i>	Return memory allocated for lock structures
<i>iofunc_lock_ocb_default()</i>	Default handler for the <i>lock_ocb</i> callout
<i>iofunc_lseek()</i>	Handle an <code>_IO_LSEEK</code> message
<i>iofunc_lseek_default()</i>	Default handler for <code>_IO_LSEEK</code> message
<i>iofunc_mknod()</i>	Verify a client's ability to make a new filesystem entry point
<i>iofunc_mmap()</i>	Handle an <code>IO_MMAP</code> message
<i>iofunc_mmap_default()</i>	Default handler for <code>IO_MMAP</code> messages
<i>iofunc_notify()</i>	Install, poll, or remove a notification handler
<i>iofunc_notify_remove()</i>	Remove notification entries from list
<i>iofunc_notify_trigger()</i>	Send notifications to queued clients
<i>iofunc_ocb_attach()</i>	Initialize an Open Control Block
<i>iofunc_ocb_alloc()</i>	Allocate an <i>iofunc</i> OCB

<i>iofunc_ocb_detach()</i>	Release OCB resources
<i>iofunc_ocb_free()</i>	Deallocate an iofunc OCBs memory
iofunc_ocb_t	Open Control Block structure
<i>iofunc_open()</i>	Verify a client's ability to open a resource
<i>iofunc_open_default()</i>	Default handler for <code>_IO_CONNECT</code> messages
<i>iofunc_openfd()</i>	Increment count and locking flags
<i>iofunc_openfd_default()</i>	Default handler for <code>_IO_OPENFD</code> messages
<i>iofunc_pathconf()</i>	Support <i>pathconf()</i> requests
<i>iofunc_pathconf_default()</i>	Default handler for <code>_IO_PATHCONF</code> messages
<i>iofunc_read_default()</i>	Default handler for <code>_IO_READ</code> messages
<i>iofunc_readlink()</i>	Verify a client's ability to read a symbolic link
<i>iofunc_read_verify()</i>	Verify a client's read access to a resource
<i>iofunc_rename()</i>	Do permission checks for a <code>_IO_CONNECT_RENAME</code> message
<i>iofunc_space_verify()</i>	Do permission checks for <code>_IO_SPACE</code> message
<i>iofunc_stat()</i>	Populate a stat structure

<i>iofunc_stat_default()</i>	Default handler for <code>_IO_STAT</code> messages
<i>iofunc_sync()</i>	Indicate if synchronization is needed
<i>iofunc_sync_default()</i>	Default handler for <code>_IO_SYNC</code> messages
<i>iofunc_sync_verify()</i>	Verify permissions to sync
<i>iofunc_time_update()</i>	Update time stamps
<i>iofunc_unblock()</i>	Unblock OCBs
<i>iofunc_unblock_default()</i>	Default unblock handler
<i>iofunc_unlink()</i>	Verify that an entry can be unlinked
<i>iofunc_unlock_ocb_default()</i>	Default handler for the <code>unlock_ocb</code> callout
<i>iofunc_utime()</i>	Update time stamps
<i>iofunc_utime_default()</i>	Default handler for <code>_IO_UTIME</code> messages
<i>iofunc_write_default()</i>	Default handler for <code>_IO_WRITE</code> messages
<i>iofunc_write_verify()</i>	Verify a client's write access to a resource
<i>ionotify()</i>	Arm a resource manager
<i>mount()</i>	Mount a filesystem

<i>mount_parse_generic_args()</i>	Strip off common mount arguments
<i>resmgr_devino()</i>	Get the device and inode number
<i>resmgr_open_bind()</i>	Associate an OCB with a process
<i>rsrdbmgr_attach()</i>	Reserve a system resource for a process
<i>rsrdbmgr_create()</i>	Create a system resource
<i>rsrdbmgr_destroy()</i>	Destroy a system resource
<i>rsrdbmgr_detach()</i>	Return a system resource to the resource database
<i>rsrdbmgr_devno_attach()</i>	Get a major and minor number
<i>rsrdbmgr_devno_detach()</i>	Detach a major and minor number
<i>rsrdbmgr_query()</i>	Query the resource database
<i>umount()</i>	Unmount a filesystem

Searching and sorting functions

These functions provide searching and sorting capabilities (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.):

<i>alphasort()</i>	Compare two directory entries
--------------------	-------------------------------

<i>bsearch()</i>	Perform a binary search on a sorted array
<i>ffs()</i>	Find the first bit set in a bit string
<i>hcreate()</i>	Create a hash search table
<i>hdestroy()</i>	Destroy the hash search table
<i>hsearch()</i>	Search the hash search table
<i>index()</i>	Find a character in a string
<i>lfind()</i>	Find entry in a linear search table
<i>lsearch()</i>	Linear search and update
<i>pathfind()</i> , <i>pathfind_r()</i>	Search for a file in a list of directories
<i>qsort()</i>	Sort an array, using a modified Quicksort algorithm
<i>re_comp()</i>	Compile a regular expression
<i>re_exec()</i>	Execute a regular expression
<i>regcomp()</i>	Compile a regular expression
<i>regerror()</i>	Explain a regular expression error code
<i>regexec()</i>	Compare a string with a compiled regular expression
<i>regfree()</i>	Release memory allocated for a regular expression
<i>rindex()</i>	Find a character in a string
<i>strcspn()</i>	Count the characters at the beginning of a string that aren't in a character set
<i>strstr()</i>	Find one string inside another

Shared memory functions

These functions provide memory mapping capabilities:

<i>mmap()</i> , <i>mmap64()</i>	Map a memory region into a process address space
<i>mprotect()</i>	Change memory protection
<i>munmap()</i>	Unmap previously mapped addresses
<i>shm_ctl()</i>	Give special attributes to a shared memory object
<i>shm_open()</i>	Open a shared memory object
<i>shm_unlink()</i>	Remove a shared memory object

Signal functions

These functions deal with handling and sending signals.

<i>DebugBreak()</i>	Enter the process debugger
<i>DebugKDBreak()</i>	Enter the kernel debugger
<i>DebugKDOutput()</i>	Print text with the kernel debugger
<i>kill()</i>	Send a signal to a process or a group of processes
<i>killpg()</i>	Send a signal to a process group
<i>pause()</i>	Suspend the process until delivery of a signal
<i>raise()</i>	Signal an exceptional condition
<i>sigaction()</i>	Examine or specify the action associated with a signal
<i>sigaddset()</i>	Add a signal to a set

<i>sigdelset()</i>	Delete a signal from a set
<i>sigemptyset()</i>	Initialize a set to contain no signals
<i>sigfillset()</i>	Initialize a set to contain all signals
<i>sigismember()</i>	See if a given signal is in a given set
<i>signal()</i>	Set handling for exceptional conditions
<i>SignalAction()</i> , <i>SignalAction_r()</i>	Examine and/or specify actions for signals
<i>SignalKill()</i> , <i>SignalKill_r()</i>	Send a signal to a process group, process, or thread
<i>SignalProcmask()</i> , <i>SignalProcmask_r()</i>	Modify or examine the signal blocked mask of a thread
<i>SignalSuspend()</i> , <i>SignalSuspend_r()</i>	Suspend a process until a signal is received
<i>SignalWaitinfo()</i> , <i>SignalWaitinfo_r()</i>	Select a pending signal
<i>sigpending()</i>	Examine the set of pending, masked signals for a process
<i>sigprocmask()</i>	Examine or change the signal mask for a process
<i>sigqueue()</i>	Queue a signal to a process
<i>sigsuspend()</i>	Replace the signal mask, and then suspend the process
<i>sigtimedwait()</i>	Wait for a signal or a timeout
<i>sigwait()</i>	Wait for a pending signal
<i>sigwaitinfo()</i>	Wait for a pending signal and get its information
<i>strsignal()</i>	Return the description of a signal

Stream I/O functions

A *stream* is the name given to a file or device that has been opened for data transmission. When a stream is opened, a pointer to a **FILE** structure is returned. This pointer is used to reference the stream when other functions are subsequently invoked.

When a program begins execution, a number of streams are already open for use:

<i>stderr</i>	Standard Error: output to the console (used for error messages)
<i>stdin</i>	Standard Input: input from the console
<i>stdout</i>	Standard Output: output to the console

You can redirect these standard streams by calling *freopen()*.

See also the section “File manipulation functions” for other functions that operate on files.

The functions in the section “Operating system I/O functions” may also be invoked (use the *fileno()* function to get the file descriptor). Since the stream functions may buffer input and output, use these functions with caution to avoid unexpected results.

(Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

<i>clearerr()</i>	Clear the end-of-file and error indicators for a stream
<i>fclose()</i>	Close a stream
<i>fcloseall()</i>	Close all open stream files, except <i>stdin</i> , <i>stdout</i> and <i>stderr</i>
<i>fdopen()</i>	Associate a stream with a file descriptor
<i>feof()</i>	Test the end-of-file indicator

<i>ferror()</i>	Test the error indicator for a stream
<i>fflush()</i>	Flush the input or output buffer for a stream
<i>fgetc()</i>	Get the next character from a file stream
<i>fgetchar()</i>	Get a character from <i>stdin</i>
<i>fgetpos()</i>	Store the current position of a stream
<i>fgets()</i>	Get a string of characters from a stream
<i>flushall()</i>	Clear all input buffers and write all output buffers
<i>fopen()</i>	Open a stream
<i>fprintf()</i>	Write output to a stream
<i>fputc()</i>	Write a character to an output stream
<i>fputchar()</i>	Write a character to <i>stdout</i>
<i>fputs()</i>	Write a character string to an output stream
<i>fread()</i>	Read elements of a given size from a stream
<i>freopen()</i>	Reopen a stream
<i>fscanf()</i>	Scan input from a stream
<i>fseek()</i> , <i>fseeko()</i>	Change the read/write position of a stream
<i>fsetpos()</i>	Set the current stream position
<i>ftell()</i> , <i>ftello()</i>	Return the current read/write position of a stream
<i>fwrite()</i>	Write a number of elements into a stream
<i>getc()</i>	Get the next character from a stream
<i>getchar()</i>	Get a character from <i>stdin</i>
<i>getchar_unlocked()</i>	Get a character from <i>stdin</i>

<i>getc_unlocked()</i>	Get the next character from a stream
<i>gets()</i>	Get a string of characters from a stream
<i>getw()</i>	Get a word from a stream
<i>perror()</i>	Print, in <i>stderr</i> , the message associated with the value of <i>errno</i>
<i>printf()</i>	Write formatted output to <i>stdout</i>
<i>putc()</i>	Write a character to an output stream
<i>putchar()</i>	Write a character to <i>stdout</i>
<i>putchar_unlocked()</i>	Write a character to <i>stdout</i>
<i>putc_unlocked()</i>	Write a character to an output stream
<i>puts()</i>	Write a string to <i>stdout</i>
<i>putw()</i>	Put a word on a stream
<i>rewind()</i>	Set the file position indicator to the beginning of the stream
<i>scanf()</i>	Scan formatted input from a stream
<i>setbuf()</i>	Associate a buffer with a stream
<i>setbuffer()</i>	Assign block buffering to a stream
<i>setlinebuf()</i>	Assign line buffering to a stream
<i>setvbuf()</i>	Associate a buffer with a stream
<i>snprintf()</i>	Write formatted output to a character array, up to a given max number of characters
<i>tmpfile()</i>	Create a temporary binary file
<i>ungetc()</i>	Push a character back onto an input stream

<i>vfprintf()</i>	Write formatted output to a file stream (varargs)
<i>vscanf()</i>	Scan input from a file stream (varargs)
<i>vprintf()</i>	Write formatted output to standard output (varargs)
<i>vscanf()</i>	Scan input from standard input (varargs)

See the section “Directory functions” for functions that are related to directories.

String manipulation functions

A *string* is an array of characters (with type `char`) that’s terminated with an extra null character (`\0`). Functions are passed only the address of the string, since the size can be determined by searching for the terminating character. (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

<i>basename()</i>	Find the part of a string after the last slash (/)
<i>bcmp()</i>	Compare a given number of characters in two strings
<i>bcopy()</i>	Copy a number of characters in one string to another
<i>fnmatch()</i>	Check to see if a file or path name matches a pattern
<i>getsubopt()</i>	Parse suboptions from a string
<i>index()</i>	Find a character in a string
<i>input_line()</i>	Get a string of characters from a file
<i>localeconv()</i>	Set numeric formatting according to the current locale
<i>re_comp()</i>	Compile a regular expression

<i>re_exec()</i>	Execute a regular expression
<i>regcomp()</i>	Compile a regular expression
<i>regerror()</i>	Explain a regular expression error code
<i>regex()</i>	Compare a string with a compiled regular expression
<i>regfree()</i>	Release memory allocated for a regular expression
<i>rindex()</i>	Find a character in a string
<i>sprintf()</i>	Print formatted output into a string
<i>scanf()</i>	Scan input from a character string
<i>straddstr()</i>	Concatenate one string on to the end of another
<i>strcasestr()</i>	Compare two strings, ignoring case
<i>strcat()</i>	Concatenate two strings
<i>strchr()</i>	Find the first occurrence of a character in a string
<i>strcmp()</i>	Compare two strings
<i>strncmp()</i>	Compare two strings, ignoring case
<i>strcoll()</i>	Compare two strings, using the locale's collating sequence
<i>strcpy()</i>	Copy a string
<i>strcspn()</i>	Count the characters at the beginning of a string that aren't in a given character set
<i>strdup()</i>	Create a duplicate of a string
<i>strerror()</i>	Map an error number to an error message
<i>stricmp()</i>	Compare two strings, ignoring case
<i>strlen()</i>	Compute the length of a string

<i>strlwr()</i>	Convert a string to lowercase
<i>strncasecmp()</i>	Compare two strings, ignoring case, up to a given length
<i>strncat()</i>	Concatenate two strings, up to a maximum length
<i>strncmp()</i>	Compare two strings, up to a given length
<i>strncpy()</i>	Copy a string, to a maximum length
<i>strnicmp()</i>	Compare two strings up to a given length, ignoring case
<i>strnset()</i>	Fill a string with a given character, to a given length
<i>strpbrk()</i>	Find the first character in a string that's in a given character set
<i>strrchr()</i>	Find the last occurrence of a character in a string
<i>strrev()</i>	Reverse a string
<i>strsep()</i>	Separate a string into pieces marked by given delimiters
<i>strset()</i>	Fill a string with a given character
<i>strspn()</i>	Count the characters at the beginning of a string that are in a given character set
<i>strstr()</i>	Find one string inside another
<i>strtok()</i>	Break a string into tokens
<i>strtok_r()</i>	Break a string into tokens (reentrant)
<i>strupr()</i>	Convert a string to uppercase
<i>strxfrm()</i>	Transform one string into another, to a given length
<i>vsprintf()</i>	Write formatted output to a buffer (varargs)

<i>vsnprintf()</i>	Write formatted output to a character array, up to a given max number of characters (varargs)
<i>vsscanf()</i>	Scan input from a string (varargs)

For related functions see these sections:

- “Conversion functions” — conversions to and from strings
- “Time functions” — formatting of dates and times
- “Memory manipulation functions” — operating on arrays without a terminating NUL character.

System database functions

The following functions are defined:

<i>crypt()</i>	Encrypt a password
<i>encrypt()</i>	Encrypt a string
<i>endgrent()</i>	Close the group database file
<i>endpwent()</i>	Close the password database file
<i>endspent()</i>	Close the shadow password database file
<i>fgetspent()</i>	Get an entry from the shadow password database
<i>getgrent()</i>	Return an entry from the group database
<i>getgrgid()</i>	Get information about the group with a given ID
<i>getgrgid_r()</i>	Get information about the group with a given ID
<i>getgrnam()</i>	Get information about the group with a given name
<i>getgrnam_r()</i>	Get information about the group with a given name
<i>getpass()</i>	Prompt for and read a password
<i>getpwent()</i>	Get an entry from the password database

<i>getpwnam()</i>	Get information about the user with a given name
<i>getpwnam_r()</i>	Get information about the user with a given name
<i>getpwuid()</i>	Get information about the user with a given ID
<i>getpwuid_r()</i>	Get information about the user with a given ID
<i>getspent(), getspent_r()</i>	Get an entry from the shadow password database
<i>getspnam(), getspnam_r()</i>	Get information about a user with a given name
<i>putspent()</i>	Put an entry into the shadow password database
<i>qnx_crypt()</i>	Encrypt a password (QNX 4)
<i>setkey()</i>	Set the key used in encryption
<i>setgrent()</i>	Rewind to the start of the group database file
<i>setpwent()</i>	Rewind the password database file
<i>setspent()</i>	Rewind the shadow password database file

System message log functions

The following functions are defined:

<i>closelog()</i>	Close the system log
<i>openlog()</i>	Open the system log
<i>setlogmask()</i>	Set the system log priority mask
<i>slogb()</i>	Send a message to the system logger
<i>slogf()</i>	Send a formatted message to the system logger
<i>slogi()</i>	Send a message to the system logger
<i>syslog()</i>	Write a message to the system log

<i>vslogf()</i>	Send a formatted message to the system logger (varargs)
<i>vsyslog()</i>	Control system log (varargs)

TCP/IP functions

These functions, prototypes and structures deal with TCP/IP network communications, database files, and the data server.

<i>accept()</i>	Accept a connection on a socket
addrinfo	TCP/IP address information
<i>bind()</i>	Bind a name to a socket
<i>bindresvport()</i>	Bind a socket to a privileged IP port
<i>connect()</i>	Initiate connection on a socket
<i>dn_comp()</i>	Compress an Internet domain name
<i>dn_expand()</i>	Expand a compressed Internet domain name
<i>ds_clear()</i>	Delete a data server variable
<i>ds_create()</i>	Create a data server variable
<i>ds_deregister()</i>	Deregister an application with the data server
<i>ds_flags()</i>	Set the flags for a data server variable
<i>ds_get()</i>	Retrieve a data server variable
<i>ds_register()</i>	Register an application with the data server
<i>ds_set()</i>	Set a data server variable
<i>endhostent()</i>	Close the TCP connection and the hosts file
<i>endnetent()</i>	Close the network database
<i>endprotoent()</i>	Close protocol name database file

<i>endservent()</i>	Close network services database file
<i>freeaddrinfo()</i>	Free an address information structure
<i>freeifaddrs()</i>	Free an address information structure
<i>gai_strerror()</i>	Return the string associated with a <i>getaddrinfo()</i> error code
<i>getaddrinfo()</i>	Get address information
<i>getdomainname()</i>	Get the domain name of the current host
<i>gethostbyaddr()</i>	Get a network host entry, given an Internet address
<i>gethostbyaddr_r()</i>	Get a network host entry, in a thread-safe manner
<i>gethostbyname()</i>	Get a network host entry, given a name
<i>gethostbyname2()</i>	Get a network host entry, given a name
<i>gethostbyname_r()</i>	Get a network host entry by name
<i>gethostent()</i>	Get the next entry from the host database
<i>gethostent_r()</i>	Get the next entry from the host database
<i>gethostname()</i>	Get the name of the current host
<i>getifaddrs()</i>	Get a network interface address
<i>getnetbyaddr()</i>	Get network entry
<i>getnetbyname()</i>	Get network entry
<i>getnetent()</i>	Get an entry from the network database
<i>getpeername()</i>	Get name of connected peer

<i>getprotobyname()</i>	Get protocol entry
<i>getprotobynumber()</i>	Get protocol entry by number
<i>getprotoent()</i>	Read next line of protocol name database file
<i>getservbyname()</i>	Get service entry
<i>getservbyport()</i>	Get service entry for a port
<i>getservent()</i>	Read the next line of network services database file
<i>getsockname()</i>	Get socket name
<i>getsockopt()</i>	Get options on socket name
<i>h_errno</i>	Host error variable
<i>herror()</i>	Print the message associated with the value of <i>h_errno</i> to standard error
hostent	Structure that describes an Internet host
<i>hstrerror()</i>	Get an error message string associated with the error return status
<i>htonl()</i>	Convert a 32-bit value from host-byte order to network-byte order
<i>htons()</i>	Convert a 16-bit value from host-byte order to network-byte order
ICMP	Internet Control Message Protocol
ICMP6	Internet Control Message Protocol for IPv6
<i>if_freenameindex()</i>	Free dynamic memory allocated by <i>if_nameindex()</i>

<i>if_indexoname()</i>	Map an interface index to its name
<i>if_nameindex()</i>	Return a list of interfaces
<i>if_nametoindex()</i>	Map an interface name to its index
<i>ifaddrs()</i>	Structure that describes an Internet host
<i>inet_addr()</i>	Convert a string into an Internet address
<i>inet_aton()</i>	Convert a string into an Internet address
<i>inet_lnaof()</i>	Convert an Internet address into a local network address
<i>inet_makeaddr()</i>	Convert a network number and a local network address into an Internet address
<i>inet_net_ntop()</i>	Convert an Internet network number to CIDR format
<i>inet_netof()</i>	Convert Internet address into a network number
<i>inet_net_pton()</i>	Convert an Internet network number from CIDR format to network format
<i>inet_network()</i>	Convert a string into an Internet network number
<i>inet_ntoa()</i>	Convert an Internet address into a string
<i>inet_ntoa_r()</i>	Convert an Internet address into a string
<i>inet_ntop()</i>	Convert a numeric network address to a string
<i>inet_pton()</i>	Convert a text host address to a numeric network address
INET6	Internet Protocol version 6 family
IP	Internet Protocol
IPsec	Internet security protocol

<i>ipsec_dump_policy()</i>	Generate readable string from IPsec policy specification
<i>ipsec_get_policylen()</i>	Get length of the IPsec policy
<i>ipsec_strerror()</i>	Error code for IPsec policy manipulation library
<i>ipsec_set_policy()</i>	Generate IPsec policy specification structure from readable string
IPv6	Internet Protocol version 6
<i>isfdtype()</i>	Determine whether a file descriptor refers to a socket
<i>listen()</i>	Listen for connections on a socket
<i>nbaconnect()</i>	Initiate a connection on a socket (nonblocking)
<i>nbaconnect_result()</i>	Get the status of the previous call to <i>nbaconnect()</i>
netent	Structure for information from the network database
<i>ntohl()</i>	Convert network-byte order value
<i>ntohs()</i>	Convert network-byte order value
protoent	Structure for information from the protocol database
<i>Raccept()</i>	Accept a connection on a socket (via a SOCKS server)
<i>Rbind()</i>	Bind a name to a socket (via a SOCKS server)
<i>rcmd()</i>	Execute a command on a remote host

<i>Rconnect()</i>	Initiate a connection on a socket (via a SOCKS server)
<i>read_main_config_file()</i>	Read the snmpd.conf file
<i>recv()</i>	Receive a message from a socket
<i>recvfrom()</i>	Receive a message from a socket
<i>recvmsg()</i>	Receive a message from a socket
<i>res_init()</i>	Initialize the Internet domain name resolver routines
<i>res_mkquery()</i>	Construct an Internet domain name query
<i>res_query()</i>	Make an Internet domain name query
<i>res_querydomain()</i>	Query the local Internet domain name server
<i>res_search()</i>	Make an Internet domain name search
<i>res_send()</i>	Send a preformatted Internet domain name query
<i>Rgetsockname()</i>	Get the name of a socket (via a SOCKS server)
<i>Rlisten()</i>	Listen for connections on a socket (via a SOCKS server)
ROUTE	System packet forwarding database
<i>Rrcmd()</i>	Execute a command on a remote host (via a SOCKS server)
<i>rresvport()</i>	Obtain a socket with a privileged address
<i>Rselect()</i>	Check for descriptors that are ready for reading or writing (via a SOCKS server)
<i>ruserok()</i>	Check the identity of a remote host

<i>send()</i>	Send a message to a socket
<i>sendmsg()</i>	Send a message to a socket
<i>sendto()</i>	Send a message to a socket
servent	Structure for information from the services database
<i>setdomainname()</i>	Set the domain name of the current host
<i>sethostent()</i>	Set the local hosts entry
<i>sethostname()</i>	Set the name of the current host
<i>setnetent()</i>	Open the network database
<i>setprotoent()</i>	Open protocol name database file
<i>setservent()</i>	Open network services database file
<i>setsockopt()</i>	Set options on socket name
<i>shutdown()</i>	Shut down part of a full-duplex connection
<i>snmp_close()</i>	Close an SNMP session
<i>snmp_free_pdu()</i>	Free an SNMP message structure
<i>snmp_open()</i>	Open an SNMP session
snmp_pdu	Structure that describes an SNMP Protocol Data Unit (transaction)
<i>snmp_pdu_create()</i>	Create an SNMP Protocol Data Unit message structure
<i>snmp_read()</i>	Read an SNMP message
<i>snmp_select_info()</i>	Get information that <i>select()</i> needs for SNMP

<i>snmp_send()</i>	Send SNMP messages
snmp_session	Structure that defines a set of transactions with similar transport characteristics
<i>snmp_timeout()</i>	Timeout during an SNMP session
<i>socketmark()</i>	Determine whether a socket is at the out-of-band mark
<i>socket()</i>	Create an endpoint for communication
<i>socketpair()</i>	Create a pair of connected sockets or a bi-directional pipe
<i>SOCKSinit()</i>	Initialize a connection with a SOCKS server
<i>sysctl()</i>	Get or set the system information
TCP	Internet Transmission Control Protocol
UDP	Internet User Datagram Protocol
UNIX	UNIX-domain protocol family

Terminal control functions

The following functions are defined:

<i>cfgetispeed()</i>	Return the input baud rate that's stored in a termios structure
<i>cfgetospeed()</i>	Return the output baud rate that's stored in a termios structure
<i>cfmakeraw()</i>	Set terminal attributes
<i>cfsetispeed()</i>	Set the input baud rate in a termios structure
<i>cfsetospeed()</i>	Set the output baud rate in a termios structure
<i>tcdrain()</i>	Wait until all output has been transmitted to a device

<i>tcdropline()</i>	Disconnect a communications line
<i>tcflow()</i>	Perform a flow-control operation on a data stream
<i>tcflush()</i>	Flush the input and/or output stream
<i>tcgetattr()</i>	Get the current terminal control settings
<i>tcgetpgrp()</i>	Get the process group ID associated with a device
<i>tcgetsid()</i>	Get the process group ID of the session leader for a controlling terminal
<i>tcgetsize()</i>	Get the size of a character device
<i>tcinject()</i>	Inject characters into a devices input buffer
<i>tcischars()</i>	Determine the number of characters waiting to be read
<i>tcsendbreak()</i>	Assert a break condition over a communications line
<i>tcsetattr()</i>	Change the terminal control settings for a device
<i>tcsetpgrp()</i>	Set the process group ID for a device
<i>tcsetsid()</i>	Make a terminal device a controlling device
<i>tcsetsize()</i>	Set the size of a character device
termios	Terminal control structure

Thread functions

These functions deal with threads and the objects used to synchronize threads:

<i>pthread_abort()</i>	Unconditionally terminate the target thread
<i>pthread_atfork()</i>	Register fork handlers

<i>pthread_attr_destroy()</i>	Destroy the thread attribute object
<i>pthread_attr_getdetachstate()</i>	Get the thread detach state attribute
<i>pthread_attr_getguardsize()</i>	Get the thread guardsize attribute
<i>pthread_attr_getinheritsched()</i>	Get the thread inherit scheduling attribute
<i>pthread_attr_getschedparam()</i>	Get the thread scheduling parameters attribute
<i>pthread_attr_getschedpolicy()</i>	Get the thread scheduling policy attribute
<i>pthread_attr_getscope()</i>	Get the thread contention scope attribute
<i>pthread_attr_getstackaddr()</i>	Get the thread stack address attribute
<i>pthread_attr_getstacklazy()</i>	Get thread stack attribute
<i>pthread_attr_getstacksize()</i>	Get the thread stack size attribute
<i>pthread_attr_init()</i>	Initialize thread attribute object
<i>pthread_attr_setdetachstate()</i>	Set the thread detach state attribute
<i>pthread_attr_setguardsize()</i>	Set the thread guardsize attribute

pthread_attr_setinheritsched()

Set the thread inherit scheduling attribute

pthread_attr_setschedparam()

Set the thread scheduling parameters attribute

pthread_attr_setschedpolicy()

Set the thread scheduling policy attribute

pthread_attr_setscope()

Set the thread contention scope attribute

pthread_attr_setstackaddr()

Set the thread stack address attribute

pthread_attr_setstacklazy()

Set thread stack attribute

pthread_attr_setstacksize()

Set the thread stack size attribute

pthread_barrierattr_destroy()

Destroy barrier attributes object

pthread_barrierattr_getpshared()

Get process-shared attribute of barrier attributes object

pthread_barrierattr_init()

Initialize barrier attributes object

pthread_barrierattr_setpshared()

Set process-shared attribute of barrier attributes object

pthread_barrier_destroy()

Destroy a barrier object

<i>pthread_barrier_init()</i>	Initialize a barrier object
<i>pthread_barrier_wait()</i>	Synchronize at a barrier
<i>pthread_cancel()</i>	Cancel thread
<i>pthread_cleanup_pop()</i>	Pop the cancellation cleanup handler
<i>pthread_cleanup_push()</i>	Push the cancellation cleanup handler
<i>pthread_condattr_destroy()</i>	Destroy the condition variable attribute object
<i>pthread_condattr_getclock()</i>	Get the clock selection condition variable attribute
<i>pthread_condattr_getpshared()</i>	Get the process-shared attribute from a condition variable attribute object
<i>pthread_condattr_init()</i>	Initialize the condition variable attribute object
<i>pthread_condattr_setclock()</i>	Set the clock selection condition variable attribute
<i>pthread_condattr_setpshared()</i>	Set the process-shared attribute in a condition variable attribute object
<i>pthread_cond_broadcast()</i>	Unblock threads waiting on a condition
<i>pthread_cond_destroy()</i>	Destroy the condition variable

<i>pthread_cond_init()</i>	Initialize the condition variable
<i>pthread_cond_signal()</i>	Unblock the thread waiting on condition variable
<i>pthread_cond_timedwait()</i>	Timed wait on the condition variable
<i>pthread_cond_wait()</i>	Wait on the condition variable
<i>pthread_create()</i>	Create a thread
<i>pthread_detach()</i>	Detach a thread from a process
<i>pthread_equal()</i>	Compare two thread IDs
<i>pthread_exit()</i>	Terminate the thread
<i>pthread_getconcurrency()</i>	Get the level of thread concurrency
<i>pthread_getcpuclockid()</i>	Return the clock ID of the CPU-time clock from a specified thread
<i>pthread_getschedparam()</i>	Get the thread scheduling parameters
<i>pthread_getspecific()</i>	Get the thread specific data value
<i>pthread_join()</i>	Join the thread
<i>pthread_key_create()</i>	Create the thread-specific data key
<i>pthread_key_delete()</i>	Delete the thread-specific data key

<i>pthread_kill()</i>	Send a signal to a thread
<i>pthread_mutexattr_destroy()</i>	Destroy the mutex attribute object
<i>pthread_mutexattr_getprioceiling()</i>	Get the priority ceiling of a mutex attribute object
<i>pthread_mutexattr_getprotocol()</i>	Get a mutex's scheduling protocol
<i>pthread_mutexattr_getpshared()</i>	Get the process-shared attribute from a mutex attribute object
<i>pthread_mutexattr_getrecursive()</i>	Get the recursive attribute from a mutex attribute object
<i>pthread_mutexattr_gettype()</i>	Get a mutex type
<i>pthread_mutexattr_init()</i>	Initialize a mutex attribute object
<i>pthread_mutexattr_setprioceiling()</i>	Set the priority ceiling of a mutex attribute object
<i>pthread_mutexattr_setprotocol()</i>	Set a mutex's scheduling protocol
<i>pthread_mutexattr_setpshared()</i>	Set the process-shared attribute in a mutex attribute object
<i>pthread_mutexattr_setrecursive()</i>	Set the recursive attribute in a mutex attribute object

pthread_mutexattr_settype()
Set a mutex type

pthread_mutex_destroy()
Destroy a mutex

pthread_mutex_getprioceiling()
Get a mutex's priority ceiling

pthread_mutex_init()
Initialize a mutex

pthread_mutex_lock()
Lock a mutex

pthread_mutex_setprioceiling()
Set a mutex's priority ceiling

pthread_mutex_timedlock()
Lock a mutex

pthread_mutex_trylock()
Attempt to lock a mutex

pthread_mutex_unlock()
Unlock a mutex

pthread_once() Dynamic package initialization

pthread_sleepon_timedwait()
Make a thread sleep while waiting

pthread_timedjoin()
Join a thread, with a time limit

pthread_rwlockattr_destroy()
Destroy a read-write lock attribute object

pthread_rwlockattr_getpshared()

Get the process-shared attribute of a read-write lock attribute object

pthread_rwlockattr_init()

Create a read-write lock attribute object

pthread_rwlockattr_setpshared()

Set the process-shared attribute of a read-write lock attribute object

pthread_rwlock_destroy()

Destroy a read/write lock

pthread_rwlock_init()

Initialize a read/write lock

pthread_rwlock_rdlock()

Acquire a shared read lock on a read/write lock

pthread_rwlock_timedrdlock()

Lock a read-write lock for writing

pthread_rwlock_timedwrlock()

Attempt to acquire an exclusive write lock on a read/write lock

pthread_rwlock_tryrdlock()

Attempt to acquire a shared read lock on a read/write lock

pthread_rwlock_trywrlock()

Attempt to acquire an exclusive write lock on a read/write lock

pthread_rwlock_unlock()

Unlock a read/write lock

<i>pthread_rwlock_wrlock()</i>	Acquire an exclusive write lock on a read/write lock
<i>pthread_self()</i>	Get the calling thread's ID
<i>pthread_setcancelstate()</i>	Set a thread's cancellation state
<i>pthread_setcanceltype()</i>	Set a thread's cancellation type
<i>pthread_setconcurrency()</i>	Set the concurrency level for a thread
<i>pthread_setschedparam()</i>	Set the thread scheduling parameters
<i>pthread_setspecific()</i>	Set a thread-specific data value
<i>pthread_sigmask()</i>	Examine and change blocked signals
<i>pthread_sleepon_broadcast()</i>	Unblock waiting threads
<i>pthread_sleepon_lock()</i>	Lock the <i>pthread_sleepon*()</i> functions
<i>pthread_sleepon_signal()</i>	Signal a sleeping thread
<i>pthread_sleepon_unlock()</i>	Unlock the <i>pthread_sleepon*()</i> functions
<i>pthread_sleepon_wait()</i>	Make a thread sleep while waiting

<i>pthread_spin_destroy()</i>	Destroy a thread spinlock
<i>pthread_spin_init()</i>	Initialize a thread spinlock
<i>pthread_spin_lock()</i>	Lock a thread spinlock
<i>pthread_spin_trylock()</i>	Try to lock a thread spinlock
<i>pthread_spin_unlock()</i>	Unlock a thread spinlock
<i>pthread_testcancel()</i>	Test the thread cancellation
<i>_sleepon_broadcast()</i>	Wake up multiple threads
<i>_sleepon_destroy()</i>	Destroy a sleepon lock
<i>_sleepon_init()</i>	Initialize a sleepon lock
<i>_sleepon_lock()</i>	Lock a sleepon lock
<i>_sleepon_signal()</i>	Wake up a single thread
<i>_sleepon_unlock()</i>	Unlock a sleepon lock
<i>_sleepon_wait()</i>	Wait on a sleepon lock

Time functions

These functions are concerned with dates and times. (Some of these functions have wide-character versions in the “Wide-character functions” section of the function summary.)

asctime(), *asctime_r()*

Convert time information to a string

clock()

Return the number of clock ticks used by the program

ClockAdjust(), *ClockAdjust_r()*

Adjust the time of a clock

ClockCycles()

Get the number of clock cycles

clock_getcpuclockid()

Return the clock ID of the CPU-time clock from a specified process

ClockId(), *ClockId_r()*

Get a clock ID for a given process and thread

ClockPeriod(), *ClockPeriod_r()*

Get or set a clock period

ClockTime(), *ClockTime_r()*

Get or set a clock

ctime(), *ctime_r()*

Convert calendar time to local time

daylight

Indicator of support for daylight saving time in the locale

difftime()

Calculate the difference between two times

ftime()

Get the current time, and store it in a structure

<i>gettimeofday()</i>	Get the current time
<i>gmtime()</i>	Convert calendar time to a broken-down time
<i>gmtime_r()</i>	Convert calendar time to a broken-down time
<i>localtime()</i>	Convert calendar time to local time
<i>localtime_r()</i>	Convert calendar time to local time
<i>mktime()</i>	Convert local time to calendar time
<i>settimeofday()</i>	Set the time and date
<i>strftime()</i>	Format a time into a string
<i>time()</i>	Determine the current calendar time
<i>TimerAlarm()</i> , <i>TimerAlarm_r()</i>	Send an alarm signal
<i>TimerCreate()</i> , <i>TimerCreate_r()</i>	Create a timer for a process
<i>TimerDestroy()</i> , <i>TimerDestroy_r()</i>	Destroy a process timer
<i>TimerInfo()</i> , <i>TimerInfo_r()</i>	Get information about a timer
<i>TimerSettime()</i> , <i>TimerSettime_r()</i>	Set the expiration time for a timer
<i>timer_timeout()</i> , <i>timer_timeout_r()</i>	Set a timeout on a blocking state
<i>TimerTimeout()</i> , <i>TimerTimeout_r()</i>	Set a timeout on a blocking state
<i>times()</i>	Get time-accounting information

<i>timezone</i>	The number of seconds by which the local time zone is earlier than UTC
tm	Structure that describes calendar time
<i>tzname</i>	The abbreviations for the time zone for standard and daylight savings time
<i>tzset()</i>	Set the time according to the current time zone

Variable-length argument list functions

Variable-length argument lists are used when a function doesn't have a fixed number of arguments. These macros provide the capability to access these arguments:

<i>va_arg()</i>	Get the next item in a list of variable arguments
<i>va_copy()</i>	Make a copy of a variable argument list
<i>va_end()</i>	Finish getting items from a variable argument list
<i>va_start()</i>	Start getting items from a variable argument list

Wide-character functions

If your application must use international characters, you'll probably need to work with Unicode and wide characters. The functions in this section are wide-character versions of many functions from the following function summary categories:

- Character manipulation functions
- Memory manipulation functions
- Stream I/O functions
- String manipulation functions
- Time functions

- Multibyte character functions
- Searching and sorting functions

The functions are:

<i>btowc()</i>	Convert a single-byte character to a wide character
<i>fgetwc()</i>	Read a wide character from a stream
<i>fgetws()</i>	Read a string of wide characters from a stream
<i>fputwc()</i>	Write a wide character to a stream
<i>fputws()</i>	Write a wide character string to an output stream
<i>fwide()</i>	Set the stream orientation
<i>fwprintf()</i>	Write wide-character output to a stream
<i>fwscanf()</i>	Scan wide-character input from a stream
<i>getwc()</i>	Read a wide character from <i>stdin</i>
<i>getwchar()</i>	Read a wide character from a stream
<i>iswalnum()</i>	Test for an alphabetic or a decimal digit wide character
<i>iswalpha()</i>	Test for an alphabetic wide character
<i>iswcntrl()</i>	Test for a control wide character
<i>iswctype()</i>	Test for an alphabetic or a decimal digit wide character
<i>iswdigit()</i>	Test for a decimal digit wide character
<i>iswgraph()</i>	Test for any graphical wide character
<i>iswlower()</i>	Test for a lowercase letter wide character
<i>iswprint()</i>	Test for a printable wide character

<i>iswpunct()</i>	Test for any punctuation wide character
<i>iswspace()</i>	Test for a whitespace wide character
<i>iswupper()</i>	Test for an uppercase wide character
<i>iswxdigit()</i>	Test for any hexadecimal digit wide character
<i>putwc()</i>	Write a wide character to a stream
<i>putwchar()</i>	Write a wide character to a stdout
<i>swprintf()</i>	Print formatted wide-character output into a string
<i>swscanf()</i>	Scan input from a wide character string
<i>towctrans()</i>	Convert a wide character in a specified manner
<i>towlower()</i>	Convert a wide character to lowercase
<i>towupper()</i>	Convert a wide character to uppercase
<i>ungetwc()</i>	Push a wide character back onto an input stream
<i>vfwprintf()</i>	Write formatted wide-character output to a file (varargs)
<i>vfwscanf()</i>	Scan input from a file (varargs)
<i>vswprintf()</i>	Write formatted wide-character output to a buffer (varargs)
<i>vswscanf()</i>	Scan input from a string (varargs)
<i>vwprintf()</i>	Write formatted wide-character output to standard output (varargs)
<i>vwscanf()</i>	Scan input from standard input (varargs)
<i>wcrtomb()</i>	Convert a wide-character code into a multibyte character (restartable)
<i>wscat()</i>	Concatenate two wide-character strings

<i>wcchr()</i>	Find the first occurrence of a wide character in a string
<i>wcscmp()</i>	Compare two wide-character strings
<i>wcscoll()</i>	Compare two wide-character strings, using the locale's collating sequence
<i>wcscopy()</i>	Copy a wide-character string
<i>wcscspn()</i>	Count the wide characters at the beginning of a string that aren't in a given character set
<i>wcsftime()</i>	Format the time into a wide-character string
<i>wcslen()</i>	Compute the length of a wide-character string
<i>wcsncat()</i>	Concatenate two wide-character strings, up to a maximum length
<i>wcsncmp()</i>	Compare two wide-character strings, up to a given length
<i>wcsncpy()</i>	Copy a wide-character string, to a maximum length
<i>wcspbrk()</i>	Find the first wide character in a string that's in a given character set
<i>wcsrtombs()</i>	Convert a wide-character string into a multibyte character string (restartable)
<i>wcsrchr()</i>	Find the last occurrence of a wide character in a string
<i>wcsspn()</i>	Count the wide characters at the beginning of a string that are in a given character set
<i>wcsstr()</i>	Find one wide-character string inside another
<i>wcstod(), wcstof(), wcstold()</i>	Convert a wide-character string into a double, float, or long double

<i>wcstoimax()</i> , <i>wcstoumax()</i>	Convert a wide-character string into an integer
<i>wcstok()</i>	Break a wide-character string into tokens
<i>wcstol()</i> , <i>wcstoll()</i>	Convert a wide-character string into a long or long long
<i>wcstombs()</i>	Convert a wide-character string into a multibyte character string
<i>wcstoul()</i> , <i>wcstoull()</i>	Convert a wide-character string into an unsigned long integer or unsigned long long
<i>wcsxfrm()</i>	Transform one wide-character string into another, to a given length
<i>wctob()</i>	Convert a wide character into a single-byte code
<i>wctomb()</i>	Convert a wide character into a multibyte character
<i>wctrans()</i>	Define a wide-character mapping
<i>wctype()</i>	Define a wide-character class
<i>wmemchr()</i>	Locate the first occurrence of a wide character in a buffer
<i>wmemcmp()</i>	Compare wide characters in two buffers
<i>wmemcpy()</i>	Copy wide characters from one buffer to another
<i>wmemmove()</i>	Copy wide characters from one buffer to another
<i>wmemset()</i>	Set wide characters in memory
<i>wprintf()</i>	Write formatted wide-character output to standard output
<i>wscanf()</i>	Scan formatted wide-character input from standard input

What's in a function description?

Each description contains the following sections:

Synopsis:

This section gives the header files that should be included within a source file that references the function or macro. It also shows an appropriate declaration for the function or for a function that could be substituted for a macro. This declaration isn't included in your program; only the header file(s) should be included.

When a pointer argument is passed to a function that doesn't modify the item indicated by that pointer, the argument is shown with `const` before the argument. For example, the following indicates that the array pointed at by *string* isn't changed:

```
const char *string
```

Arguments:

This section gives a brief description of the arguments to the function.

Library:

The section indicates the library that you need to bind with your application in order to use the function.

Description:

This section describes the function or macro.

Returns:

This section gives the return value (if any) for the function or macro.

Errors:

This section describes the special values that the function might assign to the global variable *errno*.



This section doesn't necessarily list *all* of the values that the function could set *errno* to.

See also:

This optional section provides a list of related functions or macros as well as pertinent docs to look for more information.

Examples:

This optional section gives one or more examples of the use of the function. The examples are often just code snippets, not complete programs.

Classification:

This section tells where the function or macro is commonly found, which may be helpful when porting code from one environment to another. Here are the classes:

ANSI	These functions or macros are defined by the ANSI C standard.
Large-file support	These functions support 64-bit offsets.
POSIX 1003.1	These functions are specified in the document <i>Information technology — Portable Operating System Interface (IEEE Std 1003.1, 1996 Edition)</i> . This document is being replaced by POSIX Std. 1003.1-2001.



For an up-to-date status of the many POSIX drafts/standards documents, see the PASC (Portable Applications Standards Committee of the IEEE Computer Society) report at <http://www.pasc.org/standing/sd11.html>.

POSIX 1003.1-2001

The standard incorporates the POSIX 1003.2-1992 and 1003.1-1996 standards, the approved drafts (POSIX 1003.1a, POSIX 1003.1d, POSIX 1003.1g and POSIX 1003.1j) and the Standard Unix specification. A joint technical working group — the Austin Common Standards Revision Group (CSRG) — was formed to merge these standards.

POSIX 1003.1 (Realtime Extensions)

This portion of POSIX defines optional sets of systems interfaces to support the source portability of applications with realtime requirements. Facilities include an efficient process creation mechanism, additional realtime scheduling policies, interfaces for execution time monitoring, for interacting with special devices, for improving I/O performance, and timeouts for blocking functions. The scope is to take existing realtime OS practice and add it to the base standard. The document has been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

POSIX 1003.1 (Threads)

This portion provides the POSIX base standard with interfaces and functionality to support the multiple flows of control, called *threads*, within a process. The facilities provided represent a small set of syntactic and semantic extensions to POSIX.1 in order to support a convenient interface for multithreading functions. The document has

been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

POSIX 1003.1a Some of these functions are described in the appendix to 1003.2 (*Shell and Utilities*), others are in the *System Application Program Interface (API) [C Language] — Amendment (POSIX 1003.1a Draft 15)*. The document has been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

POSIX 1003.1d (*IEEE Approved Draft: Additional Realtime Extensions*)

This portion extends the system interfaces defined by 1003.1 (Realtime Extensions). The document has been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

POSIX 1003.1g (*Draft: Protocol Independent Interfaces*)

This portion defines a programmatic interface for network process-to-process communication, such that the application may be independent of the underlying protocols. The document has been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

POSIX 1003.1j (*IEEE Approved Draft: Advanced Realtime Extensions*)

This portion extends the POSIX interfaces to provide C-language bindings for additional realtime functions for Typed Memory, Absolute Nanosleep, Barrier Synchronization, Reader/Writer Lock, Monotonic Clock, and Synchronized Clock. The document has been integrated into the IEEE POSIX Std. 1003.1-2001 spec.

QNX 4

These functions or macros are neither ANSI nor POSIX. They perform a function related to the QNX OS version 4. They may be found in other

implementations of C for personal computers with the QNX 4 OS. Use these functions with caution if portability is a consideration.



Any QNX 4 functions in the C library are provided *only* to make it easier to port QNX 4 programs. Don't use these in QNX Neutrino programs.

QNX Neutrino These functions or macros are neither ANSI nor POSIX. They perform a function related to the QNX Neutrino OS. They may be found in other implementations of C for personal computers with the QNX OS. Use these functions with caution if portability is a consideration.

SNMP Simple Network Management Protocol is a network-management protocol whose base document is *RFC 1067*. It's used to query and modify network device states.

SOCKS These functions are part of the SOCKS package consisting of a proxy server, client programs (**rftp** and **rtelnet**), and a library (**libsocks**) for adapting other applications into new client programs. For more information, see the appendix **SOCKS — A Basic Firewall**.

Unix These Unix-class functions reside on some Unix systems, but are outside of the POSIX or ANSI standards.

We've created the following Unix categories to differentiate:

Legacy Unix Functions included for backwards compatibility only. New applications shouldn't implement these functions.

Standard Unix	Functions that match XOPEN specifications. These functions are part of the IEEE POSIX Std. 1003.1-2001 spec.
Unix	Unix functions that don't fall into the above two categories. ;-)

Function safety:

This section summarizes whether or not it's safe to use the C library functions in certain situations:

Cancellation point

Indicates whether calling a function may or may not cause the thread to be terminated if a cancellation is pending.

Interrupt handler

An interrupt-safe function behaves as documented even if used in an interrupt handler. Functions flagged as interrupt-unsafe shouldn't be used in interrupt handlers.

Signal handler

A signal-safe function behaves as documented even if called from a signal handler *even if the signal interrupts a signal-unsafe function.*

Some of the signal-safe functions modify *errno* on failure. If you use any of these in a signal handler, asynchronous signals may have the side effect of modifying *errno* in an unpredictable way. If any of the code that can be interrupted checks the value of *errno* (this also applies to library calls, so you should assume that most library calls may internally check *errno*), make sure that your signal handler saves *errno* on entry and restores it on exit.

All of the above also applies to signal-unsafe functions, with one exception: if a signal handler calls a

signal-unsafe function, make sure that signal doesn't interrupt a signal-unsafe function.

Thread A thread-safe function behaves as documented even if called in a multi-threaded environment.

Most functions in the QNX Neutrino libraries are thread-safe. Even for those that aren't, there are still ways to call them safely in a multi-threaded program (e.g. by protecting the calls with a mutex). Such cases are explained in each function's description.



The “safety” designations documented in this manual are valid for the this release and could change in future versions. Floating-point functions aren't safe to use in interrupt handlers or signal handlers.

For a summary, see the Summary of Safety Information appendix.

Manifests



Manifests are used by C/C++ for compile-time changes or inspection.
Here are the defined items:

Manifest	Header file to include	Description
<code>__BEGIN_DECLS</code>	<code>sys/platform.h</code>	Denotes start of C code for a C++ compiled program.
<code>__BIGENDIAN__</code>	<code>sys/platform.h</code>	Code is compiled for a big-endian target.
<code>__CHAR_SIGNED__</code>	<code>sys/platform.h</code>	Code is compiled with the char type defaulting to signed .
<code>__CHAR_UNSIGNED__</code>	<code>sys/platform.h</code>	Code is compiled with the char type defaulting to unsigned .
<code>__END_DECLS</code>	<code>sys/platform.h</code>	Denotes end of C code for a C++ compiled program
<code>__INT_BITS__</code>	<code>sys/platform.h</code>	The number of bits in the <code>int</code> datatype.
<code>__LITTLEENDIAN__</code>	<code>sys/platform.h</code>	Code is compiled for a little-endian target.
<code>__LONG_BITS__</code>	<code>sys/platform.h</code>	The number of bits in the <code>long</code> datatype.
<code>__NTO_VERSION</code>	<code>sys/neutrino.h</code>	A version number times 100 (e.g. 2.00 is 200).
<code>__PTR_BITS__</code>	<code>sys/platform.h</code>	The number of bits in a void pointer.
<code>__OPTIMIZE__</code>	<code>sys/platform.h</code>	Code is compiled for optimization.
<code>__QNX__</code>	N/A	The target is for a QNX operating system (QNX 4 or QNX Neutrino).

continued...

Manifest	Header file to include	Description
<code>--QNXNTO--</code>	N/A	The target is the QNX Neutrino operating system.

Synopsis:

```
#include <stdlib.h>

void abort( void );
```

Library:

```
libc
```

Description:

The *abort()* function causes abnormal process termination to occur, unless the signal SIGABRT is caught and the signal handler doesn't return. The status *unsuccessful termination* is returned to the invoking process by means means of the function call *raise(SIGABRT)*.

Under QNX Neutrino, the *unsuccessful termination* status value is 6.

Returns:

The *abort()* function doesn't return to its caller.

Examples:

```
#include <stdlib.h>

int main( void )
{
    int major_error = 1;

    if( major_error )
        abort();

    /* You'll never get here. */
    return EXIT_SUCCESS;
}
```

Classification:

ANSI, POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Read the <i>Caveats</i>
Thread	Yes

Caveats:

A strictly-conforming POSIX application *can't* assume that the *abort()* function is safe to use in a signal handler on other platforms.

See also:

atexit(), *close()*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *sigaction()*, *signal()*, *spawn*()* functions, *system()*, *wait()*, *waitpid()*

Synopsis:

```
#include <stdlib.h>

int abs( int j );
```

Arguments:

j The number you want the absolute value of.

Library:

`libc`

Description:

The *abs()* function returns the absolute value of the integer argument *j*. If the result can't be represented as an `int`, a warning occurs.

Returns:

The absolute value of its argument.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    printf( "%d %d %d\n", abs (-5), abs (0), abs (5));
    return EXIT_SUCCESS;
}
```

produces the following output:

```
5 0 5
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

cabs(), *fabs()*, *labs()*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
```

```
int accept( int s,
            struct sockaddr * addr,
            socklen_t * addrlen );
```

Arguments:

s A socket that's been created with *socket()*.

addr A result parameter that's filled in with the address of the connecting entity, as known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the connection was made.

addrlen A value-result parameter. It should initially contain the amount of space pointed to by *addr*; on return it contains the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with SOCK_STREAM.

Library:

libsocket

Description:

The *accept()* function:

- 1 Extracts the first connection request on the queue of pending connections.
- 2 Creates a new socket with the same properties of *s*, where *s* is a socket that's been created with *socket()*, bound to an address with *bind()*, and is listening for connections after a *listen()*.

3 Allocates a new file descriptor for the socket.

If no pending connections are present on the queue, and the socket isn't marked as nonblocking, *accept()* blocks the caller until a connection is present. If the socket is marked as nonblocking and no pending connections are present on the queue, *accept()* returns an error as described below. The accepted socket may *not* be used to accept more connections. The original socket *s* remains open.

If you do a *select()* for read on an unconnected socket (on which a *listen()* has been done), the *select()* indicates when a connect request has occurred. In this way, an *accept()* can be made that won't block. For more information, see *select()*.

For certain protocols that require an explicit confirmation, *accept()* can be thought of as merely dequeuing the next connection request and *not* implying confirmation. Confirmation can be implied by a normal read or write on the new file descriptor, and rejection can be implied by closing the new socket.

You can obtain user-connection request data without confirming the connection by:

- Issuing a *recvmsg()* call with a *msg_iovlen* of 0 and a nonzero *msg_controllen*
Or
- Issuing a *getsockopt()* request.

Similarly, you can provide user-connection rejection information by issuing a *sendmsg()* call with only the control information, or by calling *setsockopt()*.

Returns:

A descriptor for the accepted socket, or -1 if an error occurs (*errno* is set).

Errors:

EAGAIN	Insufficient resources to create the new socket.
EBADF	Invalid descriptor <i>s</i> .
EFAULT	The <i>addr</i> parameter isn't in a writable part of the user address space.
EOPNOTSUPP	The referenced socket isn't a SOCK_STREAM socket.
ESRCH	Can't find the socket manager (<code>npm-ttcpip.so</code>).
EWOULDBLOCK	The socket is marked nonblocking and no connections are present to be accepted.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

bind(), *close()*, *connect()*, *listen()*, *select()*, *socket()*

access()

© 2004, QNX Software Systems Ltd.

Check to see if a file or directory can be accessed

Synopsis:

```
#include <unistd.h>

int access( const char * path,
           int  amode );
```

Arguments:

- path* The path to the file or directory that you want to access.
- amode* The access mode you want to check. This must be either:
- F_OK — test for file existence.
- or a bitwise ORing of the following access permissions to be checked, as defined in the header `<unistd.h>`:
- R_OK — test for read permission.
 - W_OK — test for write permission.
 - X_OK — for a directory, test for search permission. Otherwise, test for execute permission.

Library:

`libc`

Description:

The `access()` function checks to see if the file or directory specified by *path* exists and if it can be accessed with the file access permissions given by *amode*. However, unlike other functions (`open()` for example), it uses the real user ID and real group ID in place of the effective user and group IDs.

Returns:

- 0 The file or directory exists and can be accessed with the specified mode.
- 1 An error occurred (*errno* is set).

Errors:

EACCES	The permissions specified by <i>amode</i> are denied, or search permission is denied on a component of the path prefix.
EINVAL	An invalid value was specified for <i>amode</i> .
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of the <i>path</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENOENT	A component of the path isn't valid.
ENOSYS	The <i>access()</i> function isn't implemented for the filesystem specified in <i>path</i> .
ENOTDIR	A component of the path prefix isn't a directory.
EROFS	Write access was requested for a file residing on a read-only file system.

Examples:

```
#include <unistd.h>
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char **argv )
{
    if( argc!= 2 ) {
        fprintf( stderr,
            "use: readable <filename>\n" );
        return EXIT_FAILURE;
    }

    if( !access( argv[1], R_OK ) ) {
        printf( "ok to read %s\n", argv[1] );
        return EXIT_SUCCESS;
    } else {
        perror( argv[1] );
        return EXIT_FAILURE;
    }
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

chmod(), eaccess, errno, fstat(), open(), stat()

Synopsis:

```
#include <math.h>

double acos( double x );

float acosf( float x );
```

Arguments:

x The cosine for which you want to find the angle.

Library:

libm

Description:

These functions compute the arccosine (specified in radians) of *x*.

Returns:

The arccosine in the range $(0, \pi)$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", acos(.5) );

    return EXIT_SUCCESS;
}
```

produces the output:

1.047197

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

asin(), *atan()*, *atan2()*

Synopsis:

```
#include <math.h>

double acosh( double x );

float acoshf( float x );
```

Arguments:

x The value for which you want to compute the inverse hyperbolic cosine.

Library:

libm

Description:

These functions compute the inverse hyperbolic cosine (specified in radians) of *x*.

Returns:

The inverse hyperbolic cosine of *x* (specified in radians).

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", acosh( 1.5 ) );

    return EXIT_SUCCESS;
}
```

produces the output:

0.962424

Classification:

acosh() is standard Unix; *acoshf()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

asinh(), *atanh()*, *cosh()*, *errno*

Synopsis:

```
struct addrinfo {
    int ai_flags;
    int ai_family;
    int ai_socktype;
    int ai_protocol;
    size_t ai_addrlen;
    char * ai_canonname;
    struct sockaddr * ai_addr;
    struct addrinfo * ai_next
};
```

Description:

The **addrinfo** structure describes address information for use with TCP/IP. To get this information, call *getaddrinfo()*; to free a linked list of these structures, call *freeaddrinfo()*.

The **addrinfo** structure includes these members:

<i>ai_flags</i>	Flags. Includes ALPASSIVE, ALCANONNAME, and ALNUMERICHOST. For a complete list, see <code><netdb.h></code> .
<i>ai_family</i>	Protocol family. Includes PF_UNSPEC and PF_INET. For a complete list, see <code><sys/socket.h></code> .
<i>ai_socktype</i>	Socket type. Includes SOCK_STREAM and SOCK_DGRAM. For a complete list, see <code><sys/socket.h></code> .
<i>ai_protocol</i>	Protocol. Includes IPPROTO_TCP and IPPROTO_UDP. For a complete list, see <code><netinet/in.h></code> .
<i>ai_addrlen</i>	The length of the <i>ai_addr</i> member.
<i>ai_canonname</i>	The canonical name for <i>nodename</i> .

ai_addr Binary socket address.

ai_next A pointer to the next **addrinfo** structure in the linked list.

Classification:

POSIX 1003.1-2001

See also:

freeaddrinfo(), *gai_strerror()*, *getaddrinfo()*

Synopsis:

```
#include <aio.h>

int aio_cancel( int fd,
               struct aiocb * aiocbptr );
```

Library:

libc

Description:

The *aio_cancel()* function attempts to cancel one or more asynchronous I/O requests currently outstanding against a file descriptor.



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_cancel()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Synopsis:

```
#include <aio.h>

int aio_error( const struct aiocb * aiocbptr );
```

Library:

```
libc
```

Description:

The *aio_error()* function returns the error status associated with the **aiocb** structure referenced by the *aiocbptr* argument. The error status for an asynchronous I/O operation is the *errno* value that's set by the corresponding *read()*, *write()*, or *fsync()* operation.



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_error()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Synopsis:

```
#include <aio.h>

int aio_fsync( int op,
               struct aiocb * aiocbptr );
```

Library:

libc

Description:

The *aio_fsync()* function asynchronously forces all I/O operations associated with the file indicated by the file descriptor to the synchronized I/O completion state.



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_fsync()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Synopsis:

```
#include <aio.h>

int aio_read( struct aiocb * aiocbptr );
```

Library:

libc

Description:



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_read()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

aio_return()

© 2004, QNX Software Systems Ltd.

Get the return status for an asynchronous I/O operation

Synopsis:

```
#include <aio.h>

ssize_t aio_return( struct aiocb * aiocbptr );
```

Library:

libc

Description:

The *aio_return()* function returns the return status associated with the **aiocb** structure referenced by the *aiocbptr* argument. The return status for an asynchronous I/O operation is the value that's returned by the corresponding *read()*, *write()*, or *fsync()* operation.



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_return()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point No

Interrupt handler No

continued...

Safety

Signal handler	Yes
Thread	Yes

aio_suspend()

© 2004, QNX Software Systems Ltd.

Wait for asynchronous I/O operations to complete

Synopsis:

```
#include <aio.h>

int aio_suspend( const struct aiocb * const list[],
                int nent,
                const struct timespec * timeout );
```

Library:

libc

Description:

The *aio_suspend()* function suspends the calling thread until at least one of the asynchronous I/O operations referenced by the *list* argument has completed, until a signal interrupts the function, or, if *timeout* isn't NULL, until the time interval specified by *timeout* has passed.



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_suspend()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point Yes

continued...

Safety

Interrupt handler	No
Signal handler	Yes
Thread	Yes

aio_write()

© 2004, QNX Software Systems Ltd.

Asynchronously write to a file

Synopsis:

```
#include <aio.h>

int aio_write( struct aiocb * aiocbptr );
```

Library:

libc

Description:



Asynchronous I/O operations aren't currently supported.

Returns:

-1; *errno* is set.

Errors:

ENOSYS The *aio_write()* function isn't currently supported.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Synopsis:

```
#include <unistd.h>

unsigned int alarm( unsigned int seconds );
```

Arguments:

seconds The number of seconds of realtime to let elapse before raising the alarm, or zero to cancel any previous *alarm()* requests.

Library:

`libc`

Description:

The *alarm()* function causes the system to send the calling process a SIGALRM signal after a specified number of realtime seconds have elapsed. To add a handler for the signal, call *signal()* or *SignalAction()*.



Processor scheduling delays may cause the process to handle the signal after the desired time.

The *alarm()* requests aren't stacked; you can schedule only a single SIGALRM generation in this manner. If the SIGALRM hasn't yet been generated, *alarm()* reschedules the time at which the SIGALRM is generated.

Returns:

The number of seconds before the calling process is scheduled to receive a SIGALRM from the system, or zero if there was no previous *alarm()* request.

If an error occurs, an (**unsigned**) -1 is returned (*errno* is set).

Errors:

EAGAIN All timers are in use. You'll have to wait for a process to release one.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main()
{
    unsigned int  timeleft;

    printf( "Set the alarm and sleep\n" );
    alarm( 10 );
    sleep( 5 ); /* go to sleep for 5 seconds */

    /*
     * To get the time left before the SIGALRM is
     * to arrive, one must cancel the initial timer,
     * which returns the amount of time it had
     * remaining.
     */
    timeleft = alarm( 0 );
    printf( "Time left before cancel, and rearm: %d\n",
           timeleft );

    /*
     * Start a new timer that kicks us when timeleft
     * seconds have passed.
     */

    alarm( timeleft );

    /*
     * Wait until we receive the SIGALRM signal; any
     * signal kills us, though, since we don't have
     * a signal handler.
     */
    printf( "Hanging around, waiting to die\n" );
    pause();
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Requests from *alarm()*, *TimerAlarm()*, and *ualarm()* aren't "stacked;" only a single SIGALRM generator can be scheduled with these functions. If the SIGALRM signal hasn't been generated, the next call to *alarm()*, *TimerAlarm()*, or *ualarm()* reschedules it.

See also:

errno, *pause()*, *signal()*, *SignalAction()*, *sleep()*, *TimerAlarm()*, *timer_create()*, *timer_delete()*, *timer_gettime()*, *timer_settime()*, *ualarm()*

alloca()

© 2004, QNX Software Systems Ltd.

Allocate automatic space from the stack

Synopsis:

```
#include <alloca.h>

void* alloca( size_t size );
```

Arguments:

size The number of bytes of memory to allocate.

Library:

libc

Description:

The *alloca()* function allocates space for an object of *size* bytes from the stack. The allocated space is automatically discarded when the current function exits.



Don't use this function in an expression that's an argument to a function.

Returns:

A pointer to the start of the allocated memory, or NULL if an error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <string.h>
#include <malloc.h>
#include <stdlib.h>

FILE *open_err_file( char *name )
{
    char *buffer;

    /* allocate temporary buffer for file name */
    buffer = (char *)alloca( strlen( name ) + 5 );
```

```
    if( buffer ) {
        FILE *fp;

        sprintf( buffer, "%s.err", name );
        fp = fopen( buffer, "w" );

        return fp;
    }

    return (FILE *)NULL;
}

int main( void )
{
    FILE *fp;

    fp = open_err_file( "alloca_test" );
    if( fp == NULL ) {
        printf( "Unable to open error file\n" );
    } else {
        fprintf( fp, "Hello from the alloca test.\n" );
        fclose( fp );
    }

    return EXIT_SUCCESS;
}
```

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Don't use *alloca()* as an argument to a function.

See also:

calloc(), *malloc()*

Synopsis:

```
#include <sys/types.h>
#include <sys/dir.h>

struct direct {
    unsigned long d_fileno;
    unsigned short d_reclen;
    unsigned short d_namlen;
    char d_name[1];
};

int alphasort( struct direct **d1,
              struct direct **d2);
```

Arguments:

d1, d2 Pointers to the directory entries that you want to compare.

Library:

libc

Description:

The *alphasort()* function alphabetically compares two directory entries. You can use it as the *compar* argument to *scandir()*.

Returns:

- < 0 The *d1* entry precedes the *d2* entry alphabetically.
- 0 The entries are equivalent.
- > 0 The *d1* entry follows the *d2* entry alphabetically.

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

closedir(), *opendir()*, *readdir()*, *rewinddir()*, *scandir()*, *seekdir()*,
telldir()

Synopsis:

```
#include <stdlib.h>

unsigned int _amblksiz
```

Description:

The *_amblksiz* global variable holds the increment by which the “break” pointer for memory allocation is advanced when there’s no freed block large enough to satisfy a request to allocate a block of memory. You can change this at any time.

Classification:

QNX Neutrino

See also:

malloc()

argc

© 2004, QNX Software Systems Ltd.

The number of arguments passed to main()

Synopsis:

`int argc`

Description:

This global variable holds the number of arguments passed to *main()*.



This variable isn't defined in any header file. If you want to refer to it, you need to add your own **extern** statement.

Classification:

QNX Neutrino

See also:

argv, *auxv*, *getopt()*, *main()*

Synopsis:

```
char ** _argv;
```

Description:

This global variable holds a pointer to a vector containing the actual arguments passed to *main()*.



This variable isn't defined in any header file. If you want to refer to it, you need to add your own **extern** statement.

Classification:

QNX Neutrino

See also:

_argc, *_auxv*, *getopt()*, *main()*

asctime()*, *asctime_r()

© 2004, QNX Software Systems Ltd.

Convert time information to a string

Synopsis:

```
#include <time.h>

char* asctime( const struct tm* timeptr );

char* asctime_r( const struct tm* timeptr,
                 char* buf );
```

Arguments:

timeptr A pointer to a **tm** structure that contains the time that you want to convert to a string.

buf (*asctime_r()* only) A buffer in which *asctime_r()* can store the resulting string. This buffer must be large enough to hold at least 26 characters.

Library:

libc

Description:

The *asctime()* and *asctime_r()* functions convert the time information in the structure pointed to by *timeptr* into a string containing exactly 26 characters, in the form:

```
Tue May 7 10:40:27 2002\n\0
```

The *asctime()* function places the result string in a static buffer that's reused every time *asctime()* or *ctime()* is called. The result string for *asctime_r()* is contained in the buffer pointed to by *buf*.

All fields have a constant width. The newline character (`'\n'`) and a NUL character (`'\0'`) occupy the last two positions of the string.

Returns:

A pointer to the character string result, or NULL if an error occurred.

Classification:

asctime() is ANSI; *asctime_r()* is POSIX 1003.1 (Threads)

asctime()**Safety**

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	No

asctime_r()**Safety**

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *asctime()* and *ctime()* functions place their results in a static buffer that's reused for each call to *asctime()* or *ctime()*.

See also:

clock(), *ctime()*, *difftime()*, *gmtime()*, *localtime()*, *localtime_r()*, *mktime()*, *strftime()*, *time()*, **tm**, *tzset()*

asin()*, *asinf()

© 2004, QNX Software Systems Ltd.

Compute the arcsine of an angle

Synopsis:

```
#include <math.h>

double asin( double x );

float asinf( float x );
```

Arguments:

x The sine for which you want to find the angle.

Library:

libm

Description:

These functions compute the value of the arcsine (specified in radians) of *x*.

Returns:

The arcsine, in the range $(-\pi/2, \pi/2)$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", asin(.5) );

    return EXIT_SUCCESS;
}
```


produces the output:

0.523599

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acos(), *atan()*, *atan2()*, *errno*

asinh()*, *asinhf()

© 2004, QNX Software Systems Ltd.

Compute the inverse hyperbolic sine

Synopsis:

```
#include <math.h>

double asinh( double x );

float asinhf( float x );
```

Arguments:

x The value for which you want to compute the inverse hyperbolic sine.

Library:

libm

Description:

These functions compute the inverse hyperbolic sine of *x*.

Returns:

The inverse hyperbolic sine (specified in radians) of *x*.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", asinh( 0.5 ) );
    return EXIT_SUCCESS;
}
```

produces the output:

```
0.481212
```

Classification:

asinh() is standard Unix; *asinhf()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acosh(), *atanh()*, *sinh()*, *errno*

assert()

© 2004, QNX Software Systems Ltd.

Print a diagnostic message and optionally terminate the program

Synopsis:

```
#include <assert.h>

void assert( int expression );
```

Arguments:

expression Zero if you want to terminate the program; a nonzero value if you don't.

Library:

`libc`

Description:

The *assert()* macro prints a diagnostic message on the *stderr* stream, and terminates the program, using *abort()*, if *expression* is false (0).

The diagnostic message includes the *expression*, the name of the source file (the value of `_FILE_`) and the line number of the failed assertion (the value of `_LINE_`).

No action is taken if *expression* is true (nonzero).

You typically use the *assert()* macro while developing a program, to identify program logic errors. You should choose the *expression* so that it's true when the program is functioning as intended.

After the program has been debugged, you can use the special "no debug" identifier, `NDEBUG`, to remove calls to *assert()* from the program when it's recompiled. If you use the `-D` option to `gcc` or a `#define` directive to define `NDEBUG` (with any value), the C preprocessor ignores all *assert()* calls in the program source.



To remove the calls to `assert()`, you must define `NDEBUG` in the code *before* including the `<assert.h>` header file (i.e. `#include <assert.h>`).

If you define `NDEBUG`, the preprocessor also ignores the expression you pass to `assert()`. For example, if your code includes:

```
assert((fd = open("filename", O_RDWR)) != -1);
```

and you define `NDEBUG`, the preprocessor ignores the entire call to `assert()`, including the call to `open()`.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <assert.h>

void process_string( char *string )
{
    /* use assert to check argument */
    assert( string != NULL );
    assert( *string != '\0' );
    /* rest of code follows here */
}

int main( void )
{
    process_string( "hello" );
    process_string( "" );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

assert() is a macro.

See also:

abort(), *stderr*

Synopsis:

```
#include <math.h>

double atan( double x );

float atanf( float x );
```

Arguments:

x The tangent for which you want to find the angle.

Library:

`libm`

Description:

These functions compute the arctangent (specified in radians) of *x*.

Returns:

The arctangent, in the range $(-\pi/2, \pi/2)$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", atan(.5) );
    return EXIT_SUCCESS;
}
```

produces the output:

0.463648

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acos(), *asin()*, *atan2()*

Synopsis:

```
#include <math.h>

double atan2( double y,
              double x );

float atan2f( float y,
              float x );
```

Arguments:

x, y The value (y/x) for which you want to find the angle.

Library:

`libm`

Description:

These functions compute the value of the arctangent (specified in radians) of y/x , using the signs of both arguments to determine the quadrant of the return value. A domain error occurs if both arguments are zero.

Returns:

The arctangent of y/x , in the range $(-\pi, \pi)$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>
```

```
int main( void )
{
    printf( "%f\n", atan2( .5, 1. ) );

    return EXIT_SUCCESS;
}
```

produces the output:

```
0.463648
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acos(), *asin()*, *atan()*, *errno*

Synopsis:

```
#include <math.h>

double atanh( double x );

float atanhf( float x );
```

Arguments:

x The value for which you want to compute the inverse hyperbolic tangent.

Library:

libm

Description:

These functions compute the inverse hyperbolic tangent (specified in radians) of *x*.

Returns:

The inverse hyperbolic tangent of *x*.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", atanh( 0.5 ) );
    return EXIT_SUCCESS;
}
```

produces the output:

0.549306

Classification:

atanh() is standard Unix; *atanhf()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acosh(), *asinh()*, *errno*, *tanh()*

Synopsis:

```
#include <stdlib.h>

int atexit( register void (*func) (void) );
```

Arguments:

func A pointer to the function you want to be called when the program terminates normally. This function has no arguments and doesn't return a value; its prototype should be:

```
void func( void );
```

Library:

libc

Description:

The *atexit()* function registers a function to be called when the program terminates normally. If you register more than one function with *atexit()*, they're executed in a "last-in, first-out" order. Normal termination occurs either by a call to *exit()* or a return from *main()*.

You can register a total of 32 functions with *atexit()*.



The functions registered with *atexit()* aren't called when the program terminates with a call to *_exit()*.

Returns:

0 for success, or nonzero if an error occurs.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

void func1( void )
{
    printf( "last.\n" );
}

void func2( void )
{
    printf( "this " );
}

void func3( void )
{
    printf( "Do " );
}

int main( void )
{
    atexit( func1 );
    atexit( func2 );
    atexit( func3 );

    printf( "Do this first.\n" );

    return EXIT_SUCCESS;
}
```

produces the output:

```
Do this first.
Do this last.
```

Classification:

ANSI

Safety

Cancellation point No

Interrupt handler No

continued...

Safety

Signal handler	No
Thread	Yes

See also:

abort(), _exit(), exit()

atof()

© 2004, QNX Software Systems Ltd.

Convert a string into a double

Synopsis:

```
#include <stdlib.h>

double atof( const char* ptr );
```

Arguments:

ptr A pointer to the string to parse.

Library:

`libc`

Description:

The *atof()* function converts the string pointed to by *ptr* to a **double**. Calling it is equivalent to calling *strtod()* like this:

```
strtod( ptr, (char**)NULL )
```

Returns:

The converted **double**, or **0.0** if an error occurs.

Errors:

If an error occurs, *errno* is set to `ERANGE`.

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    double x;

    x = atof( "3.1415926" );
    printf( "x = %f\n", x );
    return EXIT_SUCCESS;
}
```


Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, sscanf(), strtod()

atoh()

© 2004, QNX Software Systems Ltd.

Convert a string containing a hexadecimal number into an unsigned number

Synopsis:

```
#include <stdlib.h>

unsigned atoh( const char* ptr );
```

Arguments:

ptr A pointer to the string to parse.

Library:

`libc`

Description:

The *atoh()* function converts the string pointed to by *ptr* to **unsigned** representation, assuming the string contains a hexadecimal (base 16) number.

Returns:

The converted value.

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    unsigned x;

    x = atoh( "F1A6" );
    printf( "number is %x\n", x );
    return EXIT_SUCCESS;
}
```

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

sscanf()

atoi()

© 2004, QNX Software Systems Ltd.

Convert a string into an integer

Synopsis:

```
#include <stdlib.h>

int atoi( const char* ptr );
```

Arguments:

ptr A pointer to the string to parse.

Library:

`libc`

Description:

The *atoi()* function converts the string pointed to by *ptr* to an `int`.

Returns:

The converted integer.

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    int x;

    x = atoi( "-289" );
    printf( "x = %d\n", x );
    return EXIT_SUCCESS;
}
```

produces the output:

```
x = -289
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atol(), itoa(), ltoa(), sscanf(), strtol(), strtoul(), ultoa(), utoa()

atol()*, *atoll()

© 2004, QNX Software Systems Ltd.

Convert a string into a long integer

Synopsis:

```
#include <stdlib.h>

long atol( const char* ptr );

int64_t atoll( const char* ptr );
```

Arguments:

ptr A pointer to the string to parse.

Library:

libc

Description:

The *atol()* function converts the string pointed to by *ptr* to a **long** integer; *atoll()* converts the string pointed to by *nptr* to an **int64_t** (**long long**) integer.

Returns:

atol() A **long** integer.
atoll() An **int64_t** integer.

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    long x;

    x = atol( "-289" );
    printf( "x = %d\n", x );
    return EXIT_SUCCESS;
}
```

produces the output:

x = -289

Classification:

atol() is ANSI; *atoll()* is Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atoi(), itoa(), ltoa(), sscanf(), strtol(), strtoul(), ultoa(), utoa()

atomic_add()

© 2004, QNX Software Systems Ltd.

Safely add to a variable

Synopsis:

```
#include <atomic.h>

void atomic_add( volatile unsigned * loc,
                unsigned incr );
```

Arguments:

loc A pointer to the value that you want to add to.

incr The number that you want to add.

Library:

`libc`

Description:

The *atomic_add()* function is a thread-safe way of doing an `(*loc) += incr` operation, even in a symmetric-multiprocessing system.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.

Examples:

To safely increment a counter shared between multiple threads:

```
#include <atomic.h>
...

volatile unsigned count;
...
```



```
atomic_add( &count, 1 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add_value(), *atomic_clr()*, *atomic_clr_value()*, *atomic_set()*,
atomic_set_value(), *atomic_sub()*, *atomic_sub_value()*,
atomic_toggle(), *atomic_toggle_value()*

atomic_add_value()

© 2004, QNX Software Systems Ltd.

Safely add to a variable, returning the previous value

Synopsis:

```
#include <atomic.h>

unsigned atomic_add_value( volatile unsigned * loc,
                          unsigned incr );
```

Arguments:

loc A pointer to the value that you want to add to.

incr The number that you want to add.

Library:

`libc`

Description:

The *atomic_add_value()* function is a thread-safe way of doing an `(*loc) += incr` operation, even in a symmetric-multiprocessing system.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.



The *atomic_add_value()* function may be slower than *atomic_add()*.

Returns:

The previous value of *loc*'s contents.

Examples:

To safely increment a counter shared between multiple threads:

```
#include <atomic.h>
...

volatile unsigned count;
unsigned previous;
...

previous = atomic_add_value( &count, 1 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_clr()*, *atomic_clr_value()*, *atomic_set()*,
atomic_set_value(), *atomic_sub()*, *atomic_sub_value()*,
atomic_toggle(), *atomic_toggle_value()*

atomic_clr()

© 2004, QNX Software Systems Ltd.

Safely clear a variable

Synopsis:

```
#include <atomic.h>

void atomic_clr( volatile unsigned * loc,
                unsigned bits );
```

Arguments:

loc A pointer to the value that you want to clear bits in.

bits The bits that you want to clear.

Library:

libc

Description:

The *atomic_clr()* function is a thread-safe way of doing an $(*loc) \&= \sim bits$ operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.

Examples:

To safely clear the `0x10101010` bits in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
...
```

```
atomic_clr( &flags, 0x10101010 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_set()*, *atomic_set_value()*,
atomic_sub(), *atomic_sub_value()*, *atomic_toggle()*,
atomic_toggle_value()

atomic_clr_value()

© 2004, QNX Software Systems Ltd.

Safely clear a variable, returning the previous value

Synopsis:

```
#include <atomic.h>

unsigned atomic_clr_value( volatile unsigned * loc,
                          unsigned bits );
```

Arguments:

loc A pointer to the value that you want to clear bits in.

bits The bits that you want to clear.

Library:

`libc`

Description:

The *atomic_clr_value()* function is a thread-safe way of doing an $(*loc) \&= \sim bits$ operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.



The *atomic_clr_value()* function may be slower than *atomic_clr()*.

Returns:

The previous value of *loc*'s contents.

Examples:

To safely clear the 0x10101010 bits in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
unsigned previous;
...

previous = atomic_clr_value( &flags, 0x10101010 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

*atomic_add(), atomic_add_value(), atomic_clr(), atomic_set(),
atomic_set_value(), atomic_sub(), atomic_sub_value(),
atomic_toggle(), atomic_toggle_value()*

atomic_set()

© 2004, QNX Software Systems Ltd.

Safely set bits in a variable

Synopsis:

```
#include <atomic.h>

void atomic_set( volatile unsigned * loc,
                unsigned bits );
```

Arguments:

loc A pointer to the location whose bits you want to toggle.

bits The bits that you want to set.

Library:

`libc`

Description:

The *atomic_set()* function is a thread-safe way of doing an `(*loc) |= bits` operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.

Examples:

To safely set the 1 bit in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
...
```



```
atomic_set( &flags, 0x01 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_sub(), *atomic_sub_value()*, *atomic_toggle()*,
atomic_toggle_value()

atomic_set_value()

© 2004, QNX Software Systems Ltd.

Safely set bits in a variable, returning the previous value

Synopsis:

```
#include <atomic.h>

unsigned atomic_set_value( volatile unsigned * loc,
                          unsigned bits );
```

Arguments:

loc A pointer to the location whose bits you want to toggle.
bits The bits that you want to set.

Library:

`libc`

Description:

The *atomic_set_value()* function is a thread-safe way of doing an `(*loc) |= bits` operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.



The *atomic_set_value()* function may be slower than *atomic_set()*.

Returns:

The previous value of *loc*'s contents.

Examples:

To safely set the 1 bit in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
unsigned previous;
...

previous = atomic_set_value( &flags, 0x01 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_set(), *atomic_sub()*, *atomic_sub_value()*, *atomic_toggle()*,
atomic_toggle_value()

atomic_sub()

© 2004, QNX Software Systems Ltd.

Safely subtract from a variable

Synopsis:

```
#include <atomic.h>

void atomic_sub( volatile unsigned * loc,
                unsigned decr );
```

Arguments:

loc A pointer to the value that you want to subtract from.

decr The number that you want to subtract.

Library:

libc

Description:

The *atomic_sub()* function is a thread-safe way of doing a `(*loc) -= decr` operation, even in a symmetric-multiprocessing system.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.

Examples:

Safely subtract 1 from a counter:

```
#include <atomic.h>
...

volatile unsigned count;
...
```

```
atomic_sub( &count, 1 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_set(), *atomic_set_value()*, *atomic_sub_value()*, *atomic_toggle()*,
atomic_toggle_value()

atomic_sub_value()

© 2004, QNX Software Systems Ltd.

Safely subtract from a variable, returning the previous value

Synopsis:

```
#include <atomic.h>

unsigned atomic_sub_value( volatile unsigned * loc,
                          unsigned  decr );
```

Arguments:

loc A pointer to the value that you want to subtract from.

decr The number that you want to subtract.

Library:

`libc`

Description:

The *atomic_sub_value()* function is a thread-safe way of doing a `(*loc) -= decr` operation, even in a symmetric-multiprocessing system.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.



The *atomic_sub_value()* function may be slower than *atomic_sub()*.

Returns:

The previous value of *loc*'s contents.

Examples:

Safely subtract 1 from a counter:

```
#include <atomic.h>
...

volatile unsigned count;
unsigned previous;
...

previous = atomic_sub_value( &count, 1 );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_set(), *atomic_set_value()*, *atomic_sub()*, *atomic_toggle()*,
atomic_toggle_value()

atomic_toggle()

© 2004, QNX Software Systems Ltd.

Safely toggle a variable

Synopsis:

```
#include <atomic.h>

void atomic_toggle( volatile unsigned * loc,
                   unsigned bits );
```

Arguments:

loc A pointer to the location whose bits you want to toggle.

bits The bits that you want to change.

Library:

libc

Description:

The *atomic_toggle()* function is a thread-safe way of doing an $(*loc) \hat{=} bits$ operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.

Examples:

To safely toggle the `0xdeadbeef` bits in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
...
```



```
atomic_toggle( &flags, 0xdeadbeef );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_set(), *atomic_set_value()*, *atomic_sub()*, *atomic_sub_value()*,
atomic_toggle_value()

atomic_toggle_value()

© 2004, QNX Software Systems Ltd.

Safely toggle a variable, returning the previous value

Synopsis:

```
#include <atomic.h>

unsigned atomic_toggle_value(
    volatile unsigned * loc,
    unsigned bits );
```

Arguments:

loc A pointer to the location whose bits you want to toggle.
bits The bits that you want to change.

Library:

`libc`

Description:

The *atomic_toggle_value()* function is a thread-safe way of doing an $(*loc) \hat{=} bits$ operation.

The *atomic_**() functions are guaranteed to complete without being preempted by another thread, even in a symmetric-multiprocessing system.

When modifying a variable shared between a thread and an interrupt, you *must* either disable interrupts or use the *atomic_**() functions.

The *atomic_**() functions are also useful for modifying variables that are referenced by more than one thread (that aren't necessarily in the same process) without having to use a mutex.



The *atomic_toggle_value()* function may be slower than *atomic_toggle()*.

Returns:

The previous value of *loc*'s contents.

Examples:

To safely toggle the `0xdeadbeef` bits in a flag:

```
#include <atomic.h>
...

volatile unsigned flags;
unsigned previous;
...

previous = atomic_toggle_value( &flags, 0xdeadbeef );
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

atomic_add(), *atomic_add_value()*, *atomic_clr()*, *atomic_clr_value()*,
atomic_set(), *atomic_set_value()*, *atomic_sub()*, *atomic_sub_value()*,
atomic_toggle()

auxv

© 2004, QNX Software Systems Ltd.

A pointer to a vector of auxiliary arguments to main()

Synopsis:

```
auxv_t * _auxv;
```

Description:

This global variable holds a pointer to a vector of auxiliary arguments to *main()*. For more information, see `<sys/auxv.h>`.



This variable isn't defined in any header file. If you want to refer to it, you need to add your own **extern** statement.

Classification:

QNX Neutrino

See also:

_argc, _argv, getopt(), main()

Synopsis:

```
#include <libgen.h>

char* basename( char* path );
```

Arguments:

path The string to parse.

Library:

`libc`

Description:

The *basename()* function takes the pathname pointed to by *path* and returns a pointer to the final component of the pathname, deleting any trailing “/” characters.

The *basename()* function returns:

A pointer to the string “/”

 If the string consists entirely of the “/” character

A pointer to the string “.”

 If *path* is a NULL pointer, or points to an empty string

The *basename()* function modifies the string pointed to by *path*, and returns a pointer to static storage.

Returns:

A pointer to the final component of *path*.

Examples:

```
#include <stdio.h>
#include <libgen.h>
#include <stdlib.h>

int main( int argc, char** argv )
{
    int x;

    for( x = 1; x < argc; x++ ) {
        printf( "%s\n", basename( argv[x] ) );
    }

    return EXIT_SUCCESS;
}
```

The table below shows the output of the program, given the input:

Input	Output
“/usr/lib”	“lib”
“/usr/”	“usr”
“/”	“/”

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

dirname()

bcmp()

© 2004, QNX Software Systems Ltd.

Compare a given number of characters in two strings

Synopsis:

```
#include <strings.h>

int bcmp( const void *s1,
          const void *s2,
          size_t n );
```

Arguments:

s1, s2 The strings you want to compare.

n The number of bytes to compare.

Library:

`libc`

Description:

The *bcmp()* function compares the byte string pointed to by *s1* to the string pointed to by *s2*. The number of bytes to compare is specified by *n*. NUL characters may be included in the comparison.



This function is similar to the ANSI *memcmp()* function, but tests only for equality. New code should use the ANSI function.

Returns:

0 The byte strings are identical.

1 The byte strings aren't identical.

Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <string.h>

int main( void )
{
```



```
    if( bcmp( "Hello there", "Hello world", 6 ) ) {  
        printf( "Not equal\n" );  
    } else {  
        printf( "Equal\n" );  
    }  
    return EXIT_SUCCESS;  
}
```

produces the output:

Equal

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

bcopy(), *bzero()*, *memcmp()*, *strcmp()*

bcopy()

© 2004, QNX Software Systems Ltd.

Copy a number of characters in one string to another

Synopsis:

```
#include <strings.h>

void bcopy( const void *src,
            void *dst,
            size_t n );
```

Arguments:

src The string you want to copy.

dst An existing array into which you want to copy the string.

n The number of bytes to copy.

Library:

libc

Description:

The *bcopy()* function copies the byte string pointed to by *src* (including any NUL characters) into the array pointed to by *dst*. The number of bytes to copy is specified by *n*. Copying of overlapping objects is guaranteed to work properly.



This function is similar to the ANSI *memmove()* function, but the order of arguments is different. New code should use the ANSI function.

Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <string.h>

int main( void )
{
    auto char buffer[80];
```

```
    bcopy( "Hello ", buffer, 6 );  
    bcopy( "world", &buffer[6], 6 );  
    printf( "%s\n", buffer );  
    return EXIT_SUCCESS;  
}
```

produces the output:

```
Hello world
```

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

bcmp(), *bzero()*, *memmove()*, *strcpy()*

bind()

© 2004, QNX Software Systems Ltd.

Bind a name to a socket

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>

int bind( int s,
          const struct sockaddr * name,
          socklen_t namelen );
```

Arguments:

<i>s</i>	The file descriptor to be bound.
<i>name</i>	A pointer to the sockaddr structure that holds the address to be bound to the socket. The socket length and format depend upon its address family.
<i>namelen</i>	The length of the sockaddr structure pointed to by <i>name</i> .

Library:

libsocket

Description:

When a socket is created with *socket()*, it exists in a namespace (address family) but has no name assigned to it. The *bind()* function assigns a name to that unnamed socket.



The *bind()* function assigns a local address. Use *connect()* to assign a remote address.

The rules used for binding names vary between communication domains.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EACCES The requested address is protected, and the current user has inadequate permission to access it.
- EADDRINUSE
The specified address is already in use.
- EADDRNOTAVAIL
The specified address isn't available from the local machine.
- EBADF Invalid descriptor *s*.
- EFAULT The *name* parameter isn't in a valid part of the user address space.
- EINVAL The socket is already bound to an address.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

ICMP, IP, TCP, and UDP protocols

connect(), *getsockname()*, *listen()*, *socket()*

Synopsis:

```
#include <sys/types.h>
#include <netinet/in.h>

int bindresvport( int sd,
                 struct sockaddr_in * sin );
```

Arguments:

sd The socket descriptor to bind to the port.

sin A pointer to a **sockaddr_in** structure that specifies the privileged IP port.

Library:

libsocket

Description:

The *bindresvport()* function binds a socket descriptor to a privileged IP port (i.e. a port number in the range 0-1023).



Only **root** can bind to a privileged port; this call fails for any other user.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EACCES You must be **root** to call *bindresvport()*.

EADDRINUSE

 The specified address is already in use.

EADDRNOTAVAIL

The specified address isn't available from the local machine.

EBADF Invalid descriptor *sd*.

EFAULT The *sin* parameter isn't a valid pointer to a **sockaddr_in** structure.

EINVAL The socket is already bound to a port.

EPFNOSUPPORT

The protocol family isn't supported.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	No

See also:

connect(), *getsockname()*, *listen()*, *socket()*

Synopsis:

```
#include <unistd.h>

int brk( void* endds );
```

Arguments:

endds A pointer to the new end of the data segment.

Library:

`libc`

Description:

The *brk()* function is used to change dynamically the amount of space allocated for the calling process's data segment (see the *exec** functions).

The change is made by resetting the process's break value and allocating the appropriate amount of space. The break value is the address of the first location beyond the end of the data segment. The amount of allocated space increases as the break value increases. Newly allocated space is set to zero. If, however, the same memory space is reallocated to the same process, its contents are undefined.

When a program begins execution using *execve()*, the break is set at the highest location defined by the program and data storage areas.

You can call *getrlimit()* to determine the maximum permissible size of the data segment; it isn't possible to set the break beyond the *rlim_max* value returned from *getrlimit()*, i.e:

```
end + rlim.rlim_max
```

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- ENOMEM This could mean:
- The data segment size limit, as set by *setrlimit()*, would be exceeded.
 - The maximum possible size of a data segment (compiled into the system) would be exceeded.
 - Insufficient space exists in the swap area to support the expansion.
 - Out of address space; the new break value would extend into an area of the address space defined by some previously established mapping (see *mmap()*).
- EAGAIN The total amount of system memory available for private pages is temporarily insufficient. This may occur even though the space requested was less than the maximum data segment size.

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

The behavior of *brk()* is unspecified if an application also uses any other memory functions (such as *malloc()*, *mmap()*, *free()*). The *brk()* function has been used in specialized cases where no other memory allocation function provided the same capability. Use *mmap()* instead because it can be used portably with all other memory allocation functions and with any function that uses other allocation functions.

The value of the argument to *brk()* is rounded up for alignment with eight-byte boundaries.

Setting the break may fail due to a temporary lack of swap space. It isn't possible to distinguish this from a failure caused by exceeding the maximum size of the data segment without consulting *getrlimit()*.

See also:

_btext, *_edata*, *_end*, *_etext*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *free()*, *getrlimit()*, *malloc()*, *mmap()*, *sbrk()*

bsearch()

© 2004, QNX Software Systems Ltd.

Perform a binary search on a sorted array

Synopsis:

```
#include <stdlib.h>

void *bsearch( const void *key,
               const void *base,
               size_t num,
               size_t width,
               int (*compar) ( const void *pkey,
                               const void *pbase) );
```

Arguments:

key The object to search for.

base A pointer to the first element in the array.

num The number of elements in the array.

width The size of an element, in bytes.

compare A pointer to a user-supplied function that *lfnd()* calls to compare an array element with the *key*.

The arguments to the comparison function are:

- *pkey* — the same pointer as *key*
- *pbase* — a pointer to an element in the array.

The comparison function must return an integer less than, equal to, or greater than zero if the *key* object is less than, equal to, or greater than the element in the array.

Library:

`libc`

Description:

The *bsearch()* function performs a binary search on the sorted array of *num* elements pointed to by *base*, for an item that matches the object pointed to by *key*.

Returns:

A pointer to a matching member of the array, or NULL if a matching object couldn't be found.



If there are multiple values in the array that match the *key*, the return value could be any of these duplicate values.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

static const char *keywords[] = {
    "auto",
    "break",
    "case",
    "char",
    /* ... */
    "while"
};

#define NUM_KW    sizeof(keywords) / sizeof(char *)

int kw_compare( const void *p1, const void *p2 )
{
    const char *p1c = (const char *) p1;
    const char **p2c = (const char **) p2;

    return( strcmp( p1c, *p2c ) );
}

int keyword_lookup( const char *name )
{
    const char **key;

    key = (char const **) bsearch( name, keywords,
        NUM_KW, sizeof( char * ), kw_compare );
    if( key == NULL ) return( -1 );
}
```

```
        return key - keywords;
    }

int main( void )
{
    printf( "%d\n", keyword_lookup( "case" ) );
    printf( "%d\n", keyword_lookup( "crigger" ) );
    printf( "%d\n", keyword_lookup( "auto" ) );

    return EXIT_SUCCESS;
}
```

This program produces the following output:

```
2
-1
0
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

lfind(), lsearch(), qsort()

Synopsis:

N/A

Description:

This linker symbol defines the beginning of the text segment. This variable isn't defined in any header file.

Classification:

QNX Neutrino

See also:

brk(), *_edata*, *_end*, *_etext*, *sbrk()*

btowc()

© 2004, QNX Software Systems Ltd.

Convert a single-byte character to a wide character

Synopsis:

```
#include <wchar.h>

wint_t btowc( int c );
```

Arguments:

c The single-byte character that you want to convert.

Library:

`libc`

Description:

The *btowc()* function converts the given character (if it's a valid one-byte character in the initial shift state) into a wide character.

This function is the single-byte version of *mbtowc()*.

Returns:

The wide-character representation of the character, or WEOF if *c* has the value EOF or (**unsigned char**) *c* isn't a valid one-byte character in the initial conversion state.

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

“Character manipulation functions” and “Wide-character functions”
in the summary of functions chapter.

bzero()

© 2004, QNX Software Systems Ltd.

Set the first part of an object to null bytes

Synopsis:

```
#include <strings.h>

void bzero( void *dst,
            size_t n );
```

Arguments:

dst An existing object that you want to fill with zeroes.

n The number of bytes to fill.

Library:

libc

Description:

The *bzero()* function fills the first *n* bytes of the object pointed to by *dst* with zero (NUL) bytes.



This function is similar to the ANSI *memset()* function. New code should use the ANSI function.

Examples:

```
#include <stdlib.h>
#include <string.h>

int main( void )
{
    char buffer[80];

    bzero( buffer, 80 );
    return EXIT_SUCCESS;
}
```

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

bcmp(), *bcopy()*, *memset()*, *strset()*

cabs()*, *cabsf()

© 2004, QNX Software Systems Ltd.

Compute the absolute value of a complex number

Synopsis:

```
#include <math.h>

struct __cabsargs {
    double x; /* real part */
    double y; /* imaginary part */
};

double cabs( struct __cabsargs value );

struct __cabsfargs {
    float x; /* real part */
    float y; /* imaginary part */
};

float cabsf( struct __cabsfargs value );
```

Arguments:

value The complex value that you want to get the absolute value of.

Library:

libm

Description:

These functions compute the absolute value of the complex number specified by *value*, using a calculation equivalent to:

```
sqrt( ( value.x * value.x ) + ( value.y * value.y ) );
```

Returns:

The absolute value of *value*.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

struct __cabsargs c = { -3.0, 4.0 };

int main( void )
{
    printf( "%f\n", cabs( c ) );

    return EXIT_SUCCESS;
}
```

produces the output:

5.000000

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abs(), *fabs()*, *labs()*

calloc()

© 2004, QNX Software Systems Ltd.

Allocate space for an array

Synopsis:

```
#include <stdlib.h>

void* calloc ( size_t n,
              size_t size );
```

Arguments:

n The number of array elements to allocate.

size The size, in bytes, of one array element.

Library:

`libc`

Description:

The *calloc()* function allocates space from the heap for an array of *n* objects, each of *size* bytes, and initializes them to 0.



A block of memory allocated with the *calloc()* function should be freed using the *free()* function.

Returns:

A pointer to the start of the allocated memory, or NULL if an error occurred (*errno* is set).

Errors:

ENOMEM Not enough memory.

EOK No error.

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    char* buffer;

    buffer = (char* )calloc( 80, sizeof(char) );
    if( buffer == NULL ) {
        printf( "Can't allocate memory for buffer!\n" );
        return EXIT_FAILURE;
    }

    free( buffer );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:*free(), malloc(), realloc(), sbrk()*

cbrt()*, *cbrtf()

© 2004, QNX Software Systems Ltd.

Compute the cube root of a number

Synopsis:

```
#include <math.h>

double cbrt ( double x );

float cbrtf ( float x );
```

Arguments:

x The number whose cube root you want to calculate.

Library:

libm

Description:

The *cbrt()* and *cbrtf()* functions compute the cube root of *x*.

Returns:

The cube root of *x*. If *x* is NAN, *cbrt()* returns NAN.

Examples:

```
#include <stdio.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv) {
    double a, b;

    a = 27.0;
    b = cbrt(a);
    printf("The cube root of %f is %f \n", a, b);

    return(0);
}
```

produces the output:

```
The cube root of 27.000000 is 3.000000
```


Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

sqrt()

ceil()*, *ceilf()

© 2004, QNX Software Systems Ltd.

Round up a value to the next integer

Synopsis:

```
#include <math.h>

double ceil( double x );

float ceilf( float x );
```

Arguments:

x The value you want to round.

Library:

`libm`

Description:

The *ceil()* and *ceilf()* functions round the value of *x* up to the next integer (rounding towards the “ceiling”).

Returns:

The smallest integer $\geq x$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f %f %f %f %f\n", ceil( -2.1 ),
           ceil( -2. ), ceil( 0.0 ), ceil( 2. ),
           ceil( 2.1 ) );
}
```

```
    return EXIT_SUCCESS;  
}
```

produces the output:

```
-2.000000 -2.000000 0.000000 2.000000 3.000000
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

floor()

cfmakeraw()

© 2004, QNX Software Systems Ltd.

Set terminal attributes

Synopsis:

```
#include <termios.h>

int cfmakeraw( struct termios * termios_p );
```

Arguments:

termios_p A pointer to a **termios** structure that describes the terminal's control attributes.

Library:

libc

Description:

The *cfmakeraw()* function sets the terminal attributes as follows:

```
termios_p->c_iflag &= ~(IGNBRK|BRKINT|PARMRK|ISTRIP|INLCR|IGNCR|ICRNL|IXON);
termios_p->c_oflag &= ~OPOST;
termios_p->c_lflag &= ~(ECHO|ECHONL|ICANON|ISIG|IEXTEN);
termios_p->c_cflag &= ~(CSIZE|PARENB);
termios_p->c_cflag |= CS8;
```

You can get a valid **termios** control structure for an opened device by calling *tcgetattr()*.

Returns:

0 Success.

-1 An error occurred (*errno* indicates the reason).

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *cfgetispeed()*, *cfgetospeed()*, *cfsetispeed()*, *cfsetospeed()*,
tcgetattr(), *tcsetattr()*, **termios**

cfgetispeed()

© 2004, QNX Software Systems Ltd.

*Return the input baud rate that's stored in a **termios** structure*

Synopsis:

```
#include <termios.h>

speed_t cfgetispeed(
    const struct termios* termios_p );
```

Arguments:

termios_p A pointer to a **termios** structure that describes the terminal's control attributes.

Library:

libc

Description:

The *cfgetispeed()* function returns the input baud rate that's stored in the **termios** structure pointed to by *termios_p*.

You can get a valid **termios** control structure for an opened device by calling *tcgetattr()*.

Returns:

The input baud rate stored in **termios_p*, or -1 if an error occurs (*errno* is set).

Errors:

EINVAL One of the arguments is invalid.

ENOTTY This function isn't supported by the system.

Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
```

```
int main( void )
{
    int fd;
    struct termios termios_p;
    speed_t speed;

    fd = open( "/dev/ser1", O_RDWR );
    tcgetattr( fd, &termios_p);

    /*
     * Get input baud rate
     */
    speed = cfgetispeed( &termios_p);
    printf( "Input baud: %ld\n", speed );

    close( fd );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *cfgetospeed()*, *cfsetispeed()*, *cfsetospeed()*, *tcgetattr()*, *tcsetattr()*, **termios**

cfgetospeed()

© 2004, QNX Software Systems Ltd.

*Return the output baud rate that's stored in a **termios** structure*

Synopsis:

```
#include <termios.h>

speed_t cfgetospeed(
    const struct termios* termios_p );
```

Arguments:

termios_p A pointer to a **termios** structure that describes the terminal's control attributes.

Library:

libc

Description:

The *cfgetospeed()* function returns the output baud rate that's stored in the **termios** structure pointed to by *termios_p*.

You can get a valid **termios** control structure for an opened device by calling *tcgetattr()*.

Returns:

The output baud rate stored in **termios_p*, or -1 if an error occurs (*errno* is set).

Errors:

EINVAL One of the arguments is invalid.

ENOTTY This function isn't supported by the system.

Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
```



```
int main( void )
{
    int fd;
    struct termios termios_p;
    speed_t speed;

    fd = open( "/dev/ser1", O_RDWR );
    tcgetattr( fd, &termios_p);

    /*
     * Get output baud rate
     */
    speed = cfgetospeed( &termios_p);
    printf( "Output baud: %ld\n", speed );

    close( fd );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *cfgetispeed()*, *cfsetispeed()*, *cfsetospeed()*, *tcgetattr()*, *tcsetattr()*, **termios**

cfgopen()

© 2004, QNX Software Systems Ltd.

Open a configuration file

Synopsis:

```
#include <cfgopen.h>

int cfgopen( const char * path,
             unsigned flags,
             const char * historical,
             char * namebuf,
             int nblen );
```

Arguments:

<i>path</i>	The name of the configuration file that you want to open.
<i>flags</i>	Flags that control the opening; see below.
<i>historical</i>	A optional file to open as a last resort if none of the criteria for finding the path is met. This string works like a path search order, and lets you search more than one location. You can also specify %H to substitute the hostname value into the string. Specify NULL to ignore this option.
<i>namebuf</i>	A buffer to save the pathname in. Specify NULL to ignore this option.
<i>nblen</i>	The length of the buffer pointed to by <i>namebuf</i> . Specify 0 to ignore this option.

Library:

`libc`

Description:

The *cfgopen()* function opens the configuration file named by *path*. This function is a cover function for *open()* that searches several default system locations for your files, based on specified characteristics.

The value of *flags* correspond to, and have similar limitations of, the standard *open()* flags. The *flags* value is constructed by the bitwise ORing of values from the following list, defined in the `<cfgopen.h>` header file. Applications must specify exactly one of these file-access modes in the value of *flag*:

CFGFILE_RDONLY

Open for reading only.

CFGFILE_RDWR

Open for reading and writing.

CFGFILE_WRONLY

Open for writing only.

You can also include any combination of these bits in the value of *flag*:

CFGFILE_APPEND

If set, the file offset is set to the end of the file prior to each write.

CFGFILE_CREAT

If the file doesn't exist, it's created with mode 0644, the file's user ID is set to the effective user ID of the process, and the group ID is set to the effective group ID of the process or the group ID of the file's parent directory (see *chmod()*).

CFGFILE_EXCL

If **CFGFILE_EXCL** and **CFGFILE_CREAT** are set, and the file exists, *cfgopen()* fails. The check for the existence of the file and the creation of the file if it doesn't exist is atomic with respect to other processes attempting the same operation with the same filename. Specifying **CFGFILE_EXCL** without **CFGFILE_CREAT** has no effect.

CFGFILE_TRUNC

If the file exists and is a regular file, and the file is successfully opened **CFGFILE_WRONLY** or **CFGFILE_RDWR**, the file length

is truncated to zero and the mode and owner are left unchanged. CFGFILE_TRUNC has no effect on FIFO or block or character special files or directories. Using CFGFILE_TRUNC with CFGFILE_RDONLY has no effect.

Search condition flags

In order to hint to the function where it should access or construct (in the case of CFGFILE_CREAT) *path*, there are several bits that you can specify and OR into *flags*. When specified, the bits are accessed using the following search order:

- 1 CFGFILE_USER_NODE
 \$HOME/.**cfg**/*node_name*/*path*
- 2 CFGFILE_USER
 \$HOME/.**cfg**/*path*
- 3 CFGFILE_NODE
 /**etc/host_cfg**/*node_name*/*path*
- 4 CFGFILE_GLOBAL
 path

where *node_name* is the value you get by calling *confstr()* for CS_HOSTNAME.



If the directory **/etc/host_cfg** doesn't exist on the system, the following *flags* are transformed automatically:

- CFGFILE_USER_NODE becomes CFGFILE_USER
 - CFGFILE_NODE becomes CFGFILE_GLOBAL
-

When creating a file or opening a file for writing, you can specify only one of the above location flags. Set CFGFILE_NOFD when you need only the pathname, not the file descriptor. If a directory path doesn't exist when a file is opened for creation, *cfgopen()* attempts to create the path.

Returns:

A valid file descriptor if CFGFILE_NOFD isn't specified, a nonnegative value if CFGFILE_NOFD is specified, or -1 if an error occurs.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

confstr(), *fcfgopen()*, *open()*

mib.txt, *snmpd.conf* in the *Utilities Reference*

cfree()

© 2004, QNX Software Systems Ltd.

Free allocated memory

Synopsis:

```
#include <malloc.h>

int cfree( void *ptr );
```

Arguments:

ptr A pointer to the block of memory that you want to free. It's safe to call *cfree()* with a NULL pointer.

Library:

libc

Description:

The *cfree()* function deallocates the memory block specified by *ptr*, which was previously returned by a call to *calloc()*, *malloc()* or *realloc()*.

Returns:

1

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

Calling *cfree()* on a pointer already deallocated by a call to *cfree()*, *free()*, or *realloc()* could corrupt the memory allocator's data structures.

See also:

alloca(), *calloc()*, *free()*, *malloc()*, *realloc()*, *sbrk()*

cfsetispeed()

© 2004, QNX Software Systems Ltd.

*Set the input baud rate in a **termios** structure*

Synopsis:

```
#include <termios.h>

int cfsetispeed( struct termios* termios_p,
                 speed_t speed );
```

Arguments:

termios_p A pointer to a **termios** structure that describes the terminal's control attributes.

speed The new speed. Valid values for *speed* are defined in **<termios.h>**.

Library:

libc

Description:

The *cfsetispeed()* function sets the input baud rate within the **termios** structure pointed to by *termios_p* to be *speed*.

You can get a valid **termios** control structure for an opened device by calling *tcgetattr()*.



- The new baud rate isn't effective until you call *tcsetattr()* with this modified **termios** structure.
- Attempts to set baud rates to values that aren't supported by the hardware are ignored, and cause *tcsetattr()* to return an error, but *cfsetispeed()* doesn't indicate an error.
- Attempts to set input baud rates to a value that's different from the output baud rate, when the hardware doesn't support split baud rates, cause the input baud rate to be ignored, but no error is generated.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EINVAL One of the arguments is invalid.
- ENOTTY This function isn't supported by the system.

Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main( void )
{
    int fd;
    struct termios termios_p;
    speed_t speed;

    fd = open( "/dev/ser1", O_RDWR );
    tcgetattr( fd, &termios_p);

    /*
     *   Set input baud rate
    */
}
```

```
    */
    speed = 9600;
    cfsetispeed( &termios_p, speed );
    tcsetattr( fd, TCSADRAIN, &termios_p);

    close( fd );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *cfgetispeed()*, *cfgetospeed()*, *cfsetospeed()*, *tcgetattr()*,
tcsetattr(), **termios**

Synopsis:

```
#include <termios.h>

int cfsetospeed( struct termios *termios_p,
                 speed_t speed );
```

Arguments:

termios_p A pointer to a **termios** structure that describes the terminal's control attributes.

speed The new speed. Valid values for *speed* are defined in **<termios.h>**.

Library:

libc

Description:

The *cfsetospeed()* function sets the output baud rate within the **termios** structure pointed to by *termios_p* to be *speed*.

You can get a valid **termios** control structure for an opened device by calling *tcgetattr()*.



- The new baud rate isn't effective until you call *tcsetattr()*, with this modified **termios** structure.
- Attempts to set baud rates to values that aren't supported by the hardware are ignored, and cause *tcsetattr()* to return an error, but *cfsetospeed()* doesn't indicate an error.

Setting the output baud rate to B0 causes the connection to be dropped. If *termios_p* represents a modem, the modem control lines will be turned off.

Returns:

- 0 Success.
- 1 An error occurred (*errno* indicates the reason).

Errors:

- EINVAL One of the arguments is invalid.
- ENOTTY This function isn't supported by the system.

Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main( void )
{
    int fd;
    struct termios termios_p;
    speed_t speed;

    fd = open( "/dev/ser1", O_RDWR );
    tcgetattr( fd, &termios_p);

    /*
     * Set output baud rate
     */
    speed = B9600;
    cfsetospeed( &termios_p, speed );
    tcsetattr( fd, TCSADRAIN, &termios_p);

    close( fd );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *cfgetispeed()*, *cfgetospeed()*, *cfsetispeed()*, *tcgetattr()*,
tcsetattr(), **termios**

ChannelCreate()*, *ChannelCreate_r() © 2004, QNX Software Systems

Ltd.

Create a communications channel

Synopsis:

```
#include <sys/neutrino.h>

int ChannelCreate( unsigned flags );

int ChannelCreate_r( unsigned flags );
```

Arguments:

flags Flags that can be used to request notification pulses from the kernel or request other changes in behavior; a combination of the following:

- `_NTO_CHF_COID_DISCONNECT`
- `_NTO_CHF_DISCONNECT`
- `_NTO_CHF_FIXED_PRIORITY`
- `_NTO_CHF_NET_MSG`
- `_NTO_CHF_REPLY_LEN`
- `_NTO_CHF_SENDER_LEN`
- `_NTO_CHF_THREAD_DEATH`
- `_NTO_CHF_UNBLOCK`

For more information, see below.

Library:

`libc`

Description:

The *ChannelCreate()* and *ChannelCreate_r()* kernel calls create a channel that can be used to receive messages and pulses. Once created, the channel is owned by the process and isn't bound to the creating thread.

These functions are identical, except in the way they indicate errors. See the Returns section for details.

Threads wishing to communicate with the channel attach to it by calling *ConnectAttach()*. The threads may be in the same process, or in another process on the same node (or a remote node if the network manager is running). Once attached, these threads use *MsgSendv()* or *MsgSendPulse()* to enqueue messages and pulses on the channel. Messages and pulses are enqueued in priority order.

To dequeue and read messages and pulses from a channel, use *MsgReceivev()*. Any number of threads may call *MsgReceivev()* at the same time, in which case they block and queue (if no messages or pulses are waiting) for a message or pulse to arrive. A multi-threaded I/O manager typically creates multiple threads and has them all RECEIVE-blocked on the channel.

The return value of *ChannelCreate()* is a channel ID, an `int` taken from a channel vector on the process. Most managers use a single channel for most, if not all, their communications with clients. Additional channels can be used as special channels for information.

By default, when a message is received from a channel, the thread priority of the receiver is set to match that of the thread that sent the message. This basic priority inheritance prevents priority inversion. If a message arrives at a channel and there's no thread waiting to receive it, the system boosts (if necessary) all threads in the process that have received a message from the channel in the past. This boost prevents a priority inversion of the client in the case where all threads are currently working on behalf of other clients, perhaps at a lower priority. When a thread is first created, it isn't associated with a channel until it does a *MsgReceivev()* on it. In the case of multiple channels, a thread is associated with the last channel it received from.

After receiving a message, a thread can dissociate itself from the channel by calling *MsgReceivev()* with a -1 for the channel ID. Priority inheritance can be disabled by setting `_NTO_CHF_FIXED_PRIORITY` in the *flags* argument. In this case a thread's priority isn't be affected by messages it receives on a channel.

A manager typically involves the following loop. There may be one or more threads in the loop at a time. Note that your program (not each thread) should call *ChannelCreate()* only once.

```
iov_t iov;
...
SETIOV( &iov, &msg, sizeof( msg ) );
...
chid = ChannelCreate(flags);
...
for(;;) {
    /*
     * Here's a one-part message; you could just as
     * easily receive a 20-part message by filling in the
     * iov appropriately.
     */
    rcvid = MsgReceivev(chid, &iov, 1, &info);

    /* msg is filled in by MsgReceivev() */
    switch(msg.type) {
        ...
    }

    /* iov could be filled in again to point to a new message */
    MsgReplyv(rcvid, iov, 1);
}
```

Some of the channel flags in the *flags* argument request changes from the default behavior; others request notification pulses from the kernel. The pulses are received by *MsgReceivev()* on the channel and are described by a `_pulse` structure.

The channel flags and (where appropriate) associated values for the pulse's *code* and *value* are described below.

_NTO_CHF_COID_DISCONNECT

Pulse code: `_PULSE_CODE_COIDDEATH`

Pulse value: Connection ID (*coid*) of a connection that was attached to a destroyed channel.

Deliver a pulse to this channel for each connection that belongs to the calling process when the channel that the connection is attached to is destroyed. Only one channel per process can have this flag set.



If a channel has one or both of `_NTO_CHF_COID_DISCONNECT` or `_NTO_CHF_THREAD_DEATH` set, neither flag may be set for any other channel in the process.

`_NTO_CHF_DISCONNECT`

Pulse code: `_PULSE_CODE_DISCONNECT`

Pulse value: `None`

Deliver a pulse when all connections from a process are detached (e.g. `close()`, `ConnectDetach()`, `name_close()`). If a process dies without detaching all its connections, the kernel detaches them from it. When this flag is set, the server must call `ConnectDetach(scoid)` where `scoid` is the server connection ID in the pulse message. Failure to do so leaves an invalid server connection ID that can't be reused. Over time, the server may run out of available IDs. If this flag isn't set, the kernel removes the server connection ID automatically, making it available for reuse.

`_NTO_CHF_FIXED_PRIORITY`

Suppress priority inheritance when receiving messages. Receiving threads won't change their priorities to those of the sending threads.

`_NTO_CHF_NET_MSG`

Reserved for the `io_net` resource manager.

`_NTO_CHF_REPLY_LEN`

Request that the length of the reply be included in the `dstmsglen` member of the `_msg_info` structure that `MsgReceivev()` fills in. The `dstmsglen` member is valid only if you set this channel flag when you create the channel.

Ltd.

_NTO_CHF_SENDER_LEN

Request that the length of the source message be included in the *srcmsglen* member of the `_msg_info`, structure that *MsgReceivev()* fills in. The *srcmsglen* member is valid only if you set this channel flag when you create the channel.

_NTO_CHF_THREAD_DEATH

Pulse code: `_PULSE_CODE_THREADDEATH`

Pulse value: Thread ID (*tid*)

Deliver a pulse on the death of any thread in the process that owns the channel. Only one channel per process can have this flag set.



If a channel has one or both of `_NTO_CHF_COID_DISCONNECT` or `_NTO_CHF_THREAD_DEATH` set, neither flag may be set for any other channel in the process.

_NTO_CHF_UNBLOCK

Pulse code: `_PULSE_CODE_UNBLOCK`

Pulse value: Receive ID (*rcvid*)



In most cases, you'll set the `_NTO_CHF_UNBLOCK` flag.

Deliver a pulse when a thread that's REPLY-blocked on a channel attempts to unblock before its message is replied to. This occurs between the time of a *MsgReceivev()* and a *MsgReplyv()* by the server. The sending thread may be unblocked because of a signal or a kernel timeout.

If the sending thread unblocks, *MsgReplyv()* fails. The manager may not be in a position to handle this failure. It's also possible that the client will die because of the signal and never send another message. If the manager is holding onto resources for the client (such as an

open file), it may want to receive notification that the client wants to break out of its *MsgSendv()*.

Setting the `_NTO_CHF_UNBLOCK` bit in *flags* prevents a thread that's in the REPLY-blocked state from unblocking. Instead, a pulse is sent to the channel, informing the manager that the client wishes to unblock. In the case of a signal, the signal will be pending on the client thread. When the manager replies, the client is unblocked and at that point, any pending signals are acted upon. From the client's point of view, its *MsgSendv()* will have completed normally and any signal will have arrived on the opcode following the successful kernel call.

When the manager receives the pulse, it can do one of these things:

- If it believes that it will be replying shortly, it can discard the pulse, resulting in a small latency in the unblocking, or it can signal the client. A short blocking request to a filesystem often takes this approach.
- If the reply is going to take some time or an unknown amount of time, the manager should cancel the current operation and reply back with an error or whatever data is available at this time in the reply message to the client thread. A request to a device manager waiting for input would take this approach.

Blocking states

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

ChannelCreate()

The channel ID of the newly created channel. If an error occurs, the function returns -1 and sets *errno*.

ChannelCreate_r()

The channel ID of the newly created channel. This function does **NOT** set *errno*. If an error occurs, the function returns the negative of a value from the Errors section.

Errors:

- EAGAIN All kernel channel objects are in use.
- EBUSY The `_NTO_CHF_COID_DISCONNECT` or the `_NTO_CHF_THREAD_DEATH` flag was given and another channel belonging to this process already has the same flag set.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ChannelDestroy(), *close()*, *ConnectAttach()*, *ConnectDetach()*, `_msg_info`, *MsgReceivev()*, *MsgReplyv()*, *MsgSendv()*, *MsgSendPulse()*, *name_close()*, `_pulse`

ChannelDestroy()*, *ChannelDestroy_r()

Destroy a communications channel

Synopsis:

```
#include <sys/neutrino.h>

int ChannelDestroy( int chid );

int ChannelDestroy_r( int chid );
```

Arguments:

chid The channel ID, returned by *ChannelCreate()*, of the channel that you want to destroy.

Library:

`libc`

Description:

These kernel calls remove a channel specified by the channel ID *chid* argument. Once destroyed, any attempt to receive messages or pulses on the channel will fail. Any threads that are blocked on the channel by calling *MsgReceivev()* or *MsgSendv()* will be unblocked and return with an error.

The *ChannelDestroy()* and *ChannelDestroy_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

When the channel is destroyed, all server connection IDs become invalid. The client connections are also marked invalid but remain in existence until the client removes them by calling *ConnectDetach()*. An attempt by the client to use one of these invalid connections using *MsgSendv()* or *MsgSendPulse()* will return with an error.

A server typically destroys its channels prior to its termination. If it fails to do so, the kernel destroys them automatically when the process dies.

Blocking states

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

ChannelDestroy()

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

ChannelDestroy_r()

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function may return any value in the Errors section.

Errors:

EINVAL The channel specified by *chid* doesn't exist.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

ChannelDestroy(), ChannelDestroy_r()

See also:

ChannelCreate(), MsgReceivev()

chdir()

© 2004, QNX Software Systems Ltd.

Change the current working directory

Synopsis:

```
#include <unistd.h>

int chdir( const char* path );
```

Arguments:

path The new current working directory.

Library:

`libc`

Description:

The *chdir()* function changes the current working directory to *path*, which can be relative to the current working directory or an absolute path name.

Returns:

0 Success.
-1 An error occurred (*errno* is set).

Errors:

EACCES Search permission is denied for a component of *path*.
ELOOP Too many levels of symbolic links or prefixes.
ENAMETOOLONG
 The *path* argument is longer than PATH_MAX, or a
 pathname component is longer than NAME_MAX.
ENOENT The specified *path* doesn't exist, or *path* is an empty
 string.
ENOMEM There wasn't enough memory to allocate a control
 structure.

- ENOSYS The *chdir()* function isn't implemented for the filesystem specified in *path*.
- ENOTDIR A component of *path* is not a directory.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main( int argc, char* argv[] )
{
    if( argc != 2 ) {
        fprintf( stderr, "Use: cd <directory>\n" );
        return EXIT_FAILURE;
    }

    if( chdir( argv[1] ) == 0 ) {
        printf( "Directory changed to %s\n", argv[1] );
        return EXIT_SUCCESS;
    } else {
        perror( argv[1] );
        return EXIT_FAILURE;
    }
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

There's only one current working directory per *process*. In a multithreaded application, any thread calling *chdir()* will change the current working directory for all threads in that process.

See also:

errno, *getcwd()*, *mkdir()*, *rmdir()*

Synopsis:

```
#include <sys/types.h>
#include <sys/stat.h>

int chmod( const char * path,
           mode_t mode );
```

Arguments:

path The name of the file whose permissions you want to change.

mode The new permissions for the file. For more information, see “Access permissions” in the documentation for *stat()*.

Library:

libc

Description:

The *chmod()* function changes S_ISUID, S_ISGID, S_ISVTX and the file permission bits of the file specified by the pathname pointed to by *path* to the corresponding bits in the *mode* argument. The application must ensure that the effective user ID of the process matches the owner of the file or the process has appropriate privileges to do this.

If a directory is writable and the sticky bit (S_ISVTX) is set on the directory, a process can remove or rename a file within that directory only if one or more of the following is also true:

- The effective user ID of the process matches the file’s owner ID.
- The effective user ID of the process matches the directory’s owner ID.
- The file is writable by the effective user ID of the process.
- The user is a superuser (effective user ID of 0).

If a directory has the set-group ID bit set, a file created in that directory will have the same group ID as that directory. Otherwise, the newly created file's group ID is set to the effective group ID of the creating process.

If the calling process doesn't have appropriate privileges, and if the group ID of the file doesn't match the effective group ID, and the file is a regular file, bit S_ISGID (set-group-ID on execution) in the file's mode is cleared on a successful return from *chmod()*.

If the effective user ID of the calling process is equal to the file owner, or the calling process has appropriate privileges (for example, it belongs to the superuser), *chmod()* sets S_ISUID, S_ISGID and the file permission bits, defined in the `<sys/stat.h>` header file, from the corresponding bits in the *mode* argument. These bits define access permissions for the user associated with the file, the group associated with the file and all others.

This call has no effect on file descriptors for files that are already open.

If *chmod()* succeeds, the *st_ctime* field of the file is marked for update.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EACCES Search permission is denied on a component of the path prefix.
- ELOOP Too many levels of symbolic links or prefixes.
- ENAMETOOLONG
The length of the *path* string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
- ENOTDIR A component of the path prefix isn't a directory.

ENOENT	The file doesn't exist, or the <i>path</i> arguments points to an empty string.
ENOSYS	The <i>chmod()</i> function isn't implemented for the filesystem specified in <i>path</i> .
EPERM	The effective user ID doesn't match the owner of the file, and the calling process doesn't have appropriate privileges.
EROFS	The file resides on a read-only filesystem.

Examples:

```
/*
 * Change the permissions of a list of files
 * to by read/write by the owner only
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/stat.h>

int main( int argc, char **argv )
{
    int i;
    int ecode = 0;

    for( i = 1; i < argc; i++ ) {
        if( chmod( argv[i], S_IRUSR | S_IWUSR ) == -1 ) {
            perror( argv[i] );
            ecode++;
        }
    }

    return ecode;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

chown(), errno, fchmod(), fchown(), fstat(), open(), stat()

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

int chown( const char * path,
           uid_t owner,
           gid_t group );
```

Arguments:

path The name of the file whose ownership you want to change.

owner The user ID of the new owner.

group The group ID of the new owner.

Library:

`libc`

Description:

The *chown()* function changes the user ID and group ID of the file specified by *path* to be the numeric values contained in *owner* and *group*, respectively.

If the named file is a symbolic link, *chown()* changes the ownership of the file or directory to which the symbolic link refers; *lchown()* changes the ownership of the symbolic link file itself.

Only processes with an effective user ID equal to the user ID of the file or with appropriate privileges (for example, the superuser) may change the ownership of a file.

In QNX Neutrino, the `_POSIX_CHOWN_RESTRICTED` flag (tested via the `_PC_CHOWN_RESTRICTED` flag in *pathconf()*), is enforced for *path*. This means that only the superuser may change the ownership or the group of a file to anyone. Normal users can't give a file away to another user by changing the file ownership, nor to another group by changing the group ownership.

If the *path* argument refers to a regular file, the set-user-ID (S_ISUID) and set-group-ID (S_ISGID) bits of the file mode are cleared, if the function is successful.

If *chown()* succeeds, the *st_ctime* field of the file is marked for update.

Returns:

0 Success.

-1 (no changes were made in the user ID and group ID of the file).
An error occurred (*errno* is set).

Errors:

EACCES	Search permission is denied on a component of the path prefix.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of the <i>path</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENOENT	A component of the path prefix doesn't exist, or the <i>path</i> arguments points to an empty string.
ENOSYS	The <i>chown()</i> function isn't implemented for the filesystem specified in <i>path</i> .
ENOTDIR	A component of the path prefix isn't a directory.
EPERM	The effective user ID doesn't match the owner of the file, or the calling process doesn't have appropriate privileges.
EROFS	The named file resides on a read-only filesystem.

Examples:

```
/*
 * Change the ownership of a list of files
 * to the current user/group
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main( int argc, char** argv )
{
    int i;
    int ecode = 0;

    for( i = 1; i < argc; i++ ) {
        if( chown( argv[i], getuid(), getgid() ) == -1 ) {
            perror( argv[i] );
            ecode++;
        }
    }
    exit( ecode );
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:*chmod(), errno, fchown(), fstat(), lchown(), open(), stat()*

chroot()

© 2004, QNX Software Systems Ltd.

Change the root directory

Synopsis:

```
#include <unistd.h>

int chroot( const char *path );
```

Arguments:

path The name of the new root directory.

Library:

libc

Description:

The *chroot()* function causes the *path* directory to become the root directory, the starting point for path searches for path names beginning with /. The user's working directory is unaffected.

The effective user ID of the process must be superuser to change the root directory. The .. entry in the root directory is interpreted to mean the root directory itself. Thus, .. can't be used to access files outside the subtree rooted at the root directory.

Returns:

0 Success.

-1 An error occurred; *errno* is set.

Errors:

EACCES	Search permission is denied for a component of <i>path</i> .
EBADF	The descriptor isn't valid.
EFAULT	The <i>path</i> argument points to an illegal address.
EINTR	A signal was caught during the <i>chroot()</i> function.

EIO	An I/O error occurred while reading from or writing to the filesystem.
ELOOP	Too many symbolic links were encountered in translating <i>path</i> .
EMULTIHOP	Components of <i>path</i> require hopping to multiple remote machines, and the filesystem type doesn't allow it.
ENAMETOOLONG	The length of the <i>path</i> argument exceeds {PATH_MAX} , or the length of a path component exceeds {NAME_MAX} while {_POSIX_NO_TRUNC} is in effect.
ENOENT	The named directory doesn't exist or is a null pathname.
ENOLINK	The <i>path</i> points to a remote machine and the link to that machine is no longer active.
ENOTDIR	Any component of the path name isn't a directory.
EPERM	The effective user of the calling process isn't the superuser.

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

chdir()

Synopsis:

```
#include <unistd.h>

int chsize( int filedes,
            long size );
```

Arguments:

filedes A file descriptor for the file whose size you want to change.

size The new size of the file, in bytes.

Library:

`libc`

Description:

The *chsize()* function extends or truncates the file specified by *filedes* to *size* bytes. The file is padded with NUL (`'\0'`) characters if it needs to be extended.



The *chsize()* function ignores advisory locks that may have been set with the *fcntl()* function.

Returns:

0 Success.

-1 An error occurred.

Errors:

EBADF The *filedes* argument isn't a valid file descriptor, or the file isn't opened for writing.

ENOSPC There isn't enough space left on the device to extend the file.

ENOSYS The *chsize()* function isn't implemented for the filesystem specified by *filedes*.

Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <unistd.h>
#include <fcntl.h>
#include <sys/stat.h>

int main( void )
{
    int filedes;

    filedes= open( "file", O_RDWR | O_CREAT,
                  S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP );
    if( filedes!= -1 ) {
        if( chsize( filedes, 32 * 1024L ) != 0 ) {
            printf( "Error extending file\n" );
        }
        close( filedes);

        return EXIT_SUCCESS;
    }
    return EXIT_FAILURE;
}
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

close(), creat(), errno, ftruncate(), open()

clearenv()

Clear the environment

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <stdlib.h>

int clearenv( void );
```

Library:

libc

Description:

The *clearenv()* function clears the environment area; no environment variables are defined immediately after the *clearenv()* call.

Note that *clearenv()* clears the following environment variables, which may then affect the operation of other library functions such as *spawnp()*:

- **PATH**
- **SHELL**
- **TERM**
- **TERMINFO**
- **LINES**
- **COLUMNS**
- **TZ**

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

ENOMEM Not enough memory to allocate a control structure.

Examples:

Clear the entire environment and set up a new **TZ** environment variable:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    if( clearenv() != 0 ) {
        puts( "Unable to clear the environment" );
        return EXIT_FAILURE;
    }

    setenv( "TZ", "EST5EDT", 0 );

    return EXIT_SUCCESS;
}
```

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

The *clearenv()* function manipulates the environment pointed to by the global *environ* variable.

See also:

environ, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *getenv()*, *putenv()*, *searchenv()*, *setenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*, *unsetenv()*

Synopsis:

```
#include <stdio.h>

void clearerr( FILE *fp );
```

Arguments:

fp The stream for which you want to clear the flags.

Library:

```
libc
```

Description:

The *clearerr()* function clears the end-of-file and error flags for the stream specified by *fp*.

These indicators are also cleared when the file is opened, or by an explicit call to *clearerr()* or *rewind()*.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    c = 'J';
    fp = fopen( "file", "w" );
    if( fp != NULL ) {
        fputc( c, fp );
        if( ferror( fp ) ) { /* if error      */
            clearerr( fp ); /* clear the error */
            fputc( c, fp ); /* and retry it  */
        }
    }

    fclose( fp );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

feof(), ferror(), fopen(), perror(), rewind()

Synopsis:

```
#include <time.h>

clock_t clock( void );
```

Library:

```
libc
```

Description:

The *clock()* function returns the number of clock ticks of processor time used by the program since it started executing. You can convert the number of ticks into seconds by dividing by the value `CLOCKS_PER_SEC`.



In a multithreaded program, *clock()* returns the time used by *all* threads in the application; *clock()* returns the time since the program started, not the time since a specific thread started.

Returns:

The number of clock ticks.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <time.h>
#include <stdlib.h>

void compute( void )
{
    int i, j;
    double x;

    x = 0.0;
    for( i = 1; i <= 100; i++ ) {
        for( j = 1; j <= 100; j++ ) {
            x += sqrt( (double) i * j );
        }
    }
}
```

```
    printf( "%16.7f\n", x );
}

int main( void )
{
    clock_t start_time, end_time;

    start_time = clock();
    compute();
    end_time = clock();
    printf( "Execution time was %lu seconds\n",
           (long) ((end_time - start_time) / CLOCKS_PER_SEC) );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

asctime(), asctime_r(), ctime(), difftime(), gmtime(), localtime(), localtime_r(), mktime(), strftime(), time(), tzset()

Synopsis:

```
#include <sys/neutrino.h>

int ClockAdjust( clockid_t id,
                 const struct _clockadjust * new,
                 struct _clockadjust * old );

int ClockAdjust_r( clockid_t id,
                  const struct _clockadjust * new,
                  struct _clockadjust * old );
```

Arguments:

- id* The ID of the clock you want to adjust. This must be CLOCK_REALTIME; this clock maintains the system time.
- new* NULL or a pointer to a `_clockadjust` structure that specifies how to adjust the clock. Any previous adjustment is replaced.
- The `_clockadjust` structure contains at least the following members:
- `long tick_nsec_inc` — the adjustment to be made on each clock tick, in nanoseconds.
 - `unsigned long tick_count` — the number of clock ticks over which to apply the adjustment.
- old* If not NULL, a pointer to a `_clockadjust` structure where the function can store the current adjustment (before being changed by a non-NULL *new*).

Library:

`libc`

Description:

These kernel calls let you gradually adjust the time of the clock specified by *id*. You can use these functions to speed up or slow down the system clock to synchronize it with another time source – without causing major discontinuities in the time flow.

The *ClockAdjust()* and *ClockAdjust_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

The total time adjustment, in nanoseconds, is:

```
(new->tick_count * new->tick_nsec_inc)
```

If the current clock is ahead of the desired time, you can specify a negative *tick_nsec_inc* to slow down the clock. This is preferable to setting the time backwards with the *ClockTime()* kernel call, since some programs may malfunction if time goes backwards.

Picking small values for *tick_nsec_inc* and large values for *tick_count* adjusts the time slowly, while the opposite approach adjusts it rapidly. As a rule of thumb, don't try to set a *tick_nsec_inc* that exceeds the basic clock tick as set by the *ClockPeriod()* kernel call. This would change the clock rate by more than 100% and if the adjustment is negative, it could make the clock go backwards.

You can cancel any adjustment in progress by setting *tick_count* and *tick_nsec_inc* to 0.

Superuser privileges are required to adjust the clock.

Blocking states:

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

ClockAdjust() If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

ClockAdjust_r() EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function may return any value in the Errors section.

Errors:

EFAULT	A fault occurred when the kernel tried to access the buffers provided.
EINVAL	The clock id isn't valid.
EPERM	The process tried to adjust the time without having superuser capabilities.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ClockPeriod(), *ClockTime()*

ClockCycles()

© 2004, QNX Software Systems Ltd.

Get the number of clock cycles

Synopsis:

```
#include <sys/neutrino.h>
#include <inttypes.h>

uint64_t ClockCycles( void );
```

Library:

libc

Description:

The *ClockCycles()* kernel call returns the current value of a free-running 64-bit cycle counter. This is implemented on each processor as a high-performance mechanism for timing short intervals.

Several CPU architectures have an instruction that reads such a free-running counter (e.g. x86 has the RDTSC instruction). For processors that don't implement such an instruction in hardware (e.g. a 386), the kernel emulates one. This provides a lower time resolution than if the instruction is provided (838.095345 nanoseconds on an IBM PC-compatible system).

In all cases, the `SYSPAGE_ENTRY(qtime) ->cycles_per_sec` field gives the number of *ClockCycles()* increments in one second.

Symmetric MultiProcessing systems

This function, depending on the CPU architecture, returns a value from a register that's unique to each CPU in an SMP system — for instance, the TSC (Time Stamp Counter) on an x86. These registers aren't synchronized between the CPUs. So if you call *ClockCycles()*, and then the thread migrates to another CPU and you call *ClockCycles()* again, you can't subtract the two values to get a meaningful time duration.

If you wish to use *ClockCycles()* on an SMP machine, you must use the following call to "lock" the thread to a single CPU:

ThreadCtl(_NTO_TCTL_RUNMASK, ...)

Blocking states:

This call doesn't block.

Examples:

See *SYSPAGE_ENTRY()*.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

SYSPAGE_ENTRY(), *ThreadCtl()*

clock_gettimeclockid()

© 2004, QNX Software Systems Ltd.

Return the clock ID of the CPU-time clock from a specified process

Synopsis:

```
#include <sys/types.h>
#include <time.h>

extern int clock_gettimeclockid(
    pid_t pid,
    clockid_t* clock_id );
```

Arguments:

<i>pid</i>	The process ID for the process whose clock ID you want to get.
<i>clock_id</i>	A pointer to a <code>clockid_t</code> object where the function can store the clock ID.

Library:

`libc`

Description:

The `clock_gettimeclockid()` function returns the clock ID of the CPU-time clock of the process specified by *pid*. If the process described by *pid* exists and the calling process has permission, the clock ID of this clock is stored in *clock_id*.

If *pid* is zero, the clock ID of the CPU-time clock of the process marking the call is returned in *clock_id*.

A process always has permission to obtain the CPU-time clock ID of another process.

Returns:

Zero for success, or an error value.

Errors:

ESRCH No process can be found corresponding to the specified *pid*.

Classification:

POSIX 1003.1d (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

clock_getres(), *clock_gettime()*, *ClockId()*, *clock_settime()*,
pthread_getcpuclockid(), *timer_create()*

clock_getres()

© 2004, QNX Software Systems Ltd.

Get the resolution of the clock

Synopsis:

```
#include <time.h>

int clock_getres( clockid_t clock_id,
                 struct timespec * res );
```

Arguments:

clock_id The ID of the clock whose resolution you want to get.

res A pointer to a `timespec` structure in which `clock_getres()` can store the resolution. The function sets the `tv_sec` member to 0, and the `tv_nsec` member to be the resolution of the clock, in nanoseconds.

Library:

`libc`

Description:

The `clock_getres()` function gets the resolution of the clock specified by `clock_id` and puts it into the buffer pointed to by `res`.

Returns:

0 Success

-1 An error occurred (`errno` is set).

Errors:

EFAULT A fault occurred trying to access the buffers provided.

EINVAL Invalid `clock_id`.

Examples:

```
/*
 * This program prints out the clock resolution.
 */
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

int main( void )
{
    struct timespec res;

    if ( clock_getres( CLOCK_REALTIME, &res) == -1 ) {
        perror( "clock get resolution" );
        return EXIT_FAILURE;
    }
    printf( "Resolution is %ld micro seconds.\n",
           res.tv_nsec / 1000 );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:*clock_gettime()*, *clock_settime()*, *ClockPeriod()*, *timespec*

clock_gettime()

© 2004, QNX Software Systems Ltd.

Get the current time of a clock

Synopsis:

```
#include <time.h>

int clock_gettime( clockid_t clock_id,
                  struct timespec * tp );
```

Arguments:

- clock_id* The ID of the clock whose time you want to get.
- tp* A pointer to a `timespec` structure where `clock_gettime()` can store the time. This function sets the members as follows:
- *tv_sec* — the number of seconds since 1970.
 - *tv_nsec* — the number of nanoseconds expired in the current second. This value increases by some multiple of nanoseconds, based on the system clock's resolution.

Library:

`libc`

Description:

The `clock_gettime()` function gets the current time of the clock specified by *clock_id*, and puts it into the buffer pointed to by *tp*.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EFAULT A fault occurred trying to access the buffers provided.
- EINVAL Invalid *clock_id*.
- ESRCH The process associated with this request doesn't exist.

Examples:

```
/*
 * This program calculates the time required to
 * execute the program specified as its first argument.
 * The time is printed in seconds, on standard out.
 */
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>
#include <time.h>

#define BILLION 1000000000L;

int main( int argc, char** argv )
{
    struct timespec start, stop;
    double accum;

    if( clock_gettime( CLOCK_REALTIME, &start) == -1 ) {
        perror( "clock_gettime" );
        return EXIT_FAILURE;
    }

    system( argv[1] );

    if( clock_gettime( CLOCK_REALTIME, &stop) == -1 ) {
        perror( "clock_gettime" );
        return EXIT_FAILURE;
    }

    accum = ( stop.tv_sec - start.tv_sec )
            + (double)( stop.tv_nsec - start.tv_nsec )
              / (double)BILLION;
    printf( "%lf\n", accum );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

clock_getres(), *clock_settime()*, *errno*, **timespec**

Synopsis:

```
#include <time.h>

int clock_nanosleep( clockid_t clock_id,
                    int flags,
                    const struct timespec * rntp,
                    struct timespec * rntp );
```

Arguments:

clock_id The ID of the clock to use to measure the time. The possible clock types are:

CLOCK_MONOTONIC

A clock that always increases at a constant rate and can't be adjusted.

CLOCK_SOFTTIME

Same as CLOCK_REALTIME, but if the CPU is in powerdown mode, the clock stops running.

CLOCK_REALTIME

A clock that maintains the system time.

The *clock_nanosleep()* function fails if the *clock_id* argument refers to the CPU-time clock of the calling thread.

flags Flags that specify when the current thread is to be suspended from execution:

- when the time interval specified by the *rntp* argument has elapsed (TIMER_ABSTIME is *not* set).
- when the time value of the clock specified by *clock_id* reaches the absolute time specified by the *rntp* argument (TIMER_ABSTIME is set).

If, at the time of the call, the time value specified by *rntp* is less than or equal to the time value of the

specified clock, then *clock_nanosleep()* returns immediately, and the calling process isn't suspended.

- when a signal is delivered to the calling thread, and the signal's action is to invoke a signal-catching function or terminate the process.

Calling *clock_nanosleep()* with `TIMER_ABSTIME` not set, and *clock_id* set to `CLOCK_REALTIME` is the equivalent to calling *nanosleep()* with the same *rntp* and *rmtp* arguments.

rntp A pointer to a `timespec` structure that specifies the time interval between the requested time and the time actually slept.

rmtp NULL, or a pointer to a `timespec` in which the function can store the amount of time remaining in an interval.

For the relative *clock_nanosleep()* function, if *rmtp* isn't NULL, the `timespec` structure referenced by it is updated to contain the amount of time remaining in the interval (the requested time minus the time actually slept). If it's NULL, the remaining time isn't returned.

The absolute *clock_nanosleep()* function has no effect on the structure referenced by *rmtp*.

Library:

`libc`

Description:

The *clock_nanosleep()* function suspends the current thread from execution until:

- If `TIMER_ABSTIME` is set, the time value of the clock specified by *clock_id* reaches the absolute time specified by the *rntp* argument.

Or:

- If `TIMER_ABSTIME` is *not* set, the time interval specified by the *rqtp* argument has elapsed.
- Or:
- A signal is delivered to the calling thread, and the signal's action is to invoke a signal-catching function or terminate the process.

The *nanosleep()* function always uses `CLOCK_REALTIME`.

The suspension time may be longer than requested because the argument value is rounded up to an integer multiple of the sleep resolution, or because of scheduling and other system activity. Except for the case of being interrupted by a signal, the suspension time for:

- the relative *clock_nanosleep()* function (`TIMER_ABSTIME` not set) — isn't less than the time interval specified by *rqtp*, as measured by the corresponding clock
- the absolute *clock_nanosleep()* function (`TIMER_ABSTIME` set) — is in effect at least until the value of the corresponding clock reaches the absolute time specified by *rqtp*, except for the case of being interrupted by a signal.

Using the *clock_nanosleep()* function has no effect on the action or blockage of any signal.

Returns:

Zero if the requested time has elapsed, or a corresponding error value if *clock_nanosleep()* has been interrupted by a signal, or fails.

Errors:

<code>EINTR</code>	The call was interrupted by a signal.
<code>EINVAL</code>	The <i>rqtp</i> argument specified a nanosecond value less than zero or greater than or equal to 1000 million; or <code>TIMER_ABSTIME</code> is specified in <i>flags</i> and the <i>rqtp</i> argument is outside the range for the clock specified by <i>clock_id</i> ; or the <i>clock_id</i> argument doesn't specify a

known clock, or specifies the CPU-time clock of the calling thread.

ENOTSUP The *clock_id* argument specifies a clock for which *clock_nanosleep()* isn't supported, such as a CPU-time clock.

Classification:

POSIX 1003.1j (draft)

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

clock_settime(), *nanosleep()*, *sleep()*, **timespec**

Synopsis:

```
#include <time.h>

int clock_settime( clockid_t clock_id,
                  const struct timespec * tp );
```

Arguments:

- clock_id* The ID of the clock you want to set.
- tp* A pointer to a `timespec` structure containing at least the following members:
- *tv_sec* — the number of seconds since 1970.
 - *tv_nsec* — the number of nanoseconds in the current second. This value increases by some multiple of nanoseconds, based on the system clock's resolution.

Library:

`libc`

Description:

The `clock_settime()` function sets the clock specified by *clock_id*, from the buffer pointed to by *tp*.



Be careful if you set the date during the period that a time zone is switching daylight saving time (DST) to standard time. When a time zone changes to standard time, the local time goes back one hour (for example, 2:00 a.m. becomes 1:00 a.m.). The local time during this hour is ambiguous (e.g. 1:14 a.m. occurs twice in the morning that the time zone switches to standard time). To avoid problems, use UTC time to set the date in this period.

Returns:

- 0 Success
- 1 An error occurred (*errno* is set).

Errors:

- EINVAL Invalid *clock_id* or the number of nanoseconds specified by the *tv_nsec* is less than zero or greater than or equal to 1000 million.
- EPERM You don't have sufficient privilege to change the time.

Examples:

```
/* This program sets the clock forward 1 day. */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <time.h>

int main( void )
{
    struct timespec stime;

    if( clock_gettime( CLOCK_REALTIME, &stime) == -1 ) {
        perror( "getclock" );
        return EXIT_FAILURE;
    }

    stime.tv_sec += (60*60)*24L; /* Add one day */
    stime.tv_nsec = 0;
    if( clock_settime( CLOCK_REALTIME, &stime) == -1 ) {
        perror( "setclock" );
        return EXIT_FAILURE;
    }
    return EXIT_SUCCESS;
}
```


Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

clock_getres(), *clock_gettime()*, *errno*, *timespec*

ClockId()*, *ClockId_r()

© 2004, QNX Software Systems Ltd.

Get a clock ID for a given process and thread

Synopsis:

```
#include <sys/neutrino.h>
#include <inttypes.h>

extern int ClockId( pid_t pid,
                   int tid );

extern int ClockId_r( pid_t pid,
                     int tid );
```

Arguments:

- pid* The ID of the process that you want to calculate the execution time for. If this argument is zero, the ID of the process making the call is assumed.
- tid* The ID of the thread that you want to calculate the execution time for, or 0 to get the execution time for the process as a whole.

Library:

`libc`

Description:

The *ClockId()* and *ClockId_r()* kernel calls return an integer that you can pass as a `clockid_t` to *ClockTime()*. When you pass this clock ID to *ClockTime()*, the function returns (in the location pointed to by *old*) the number of nanoseconds that the specified thread of the specified process has executed.

The *ClockId()* and *ClockId_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

If the *tid* is zero, the number of nanoseconds that the process as a whole has executed is returned. On an SMP box, this number may exceed the realtime number of nanoseconds that have elapsed because multiple threads in the process can run on several CPUs at the same time.

Blocking states:

This call doesn't block.

Returns:

ClockId() An integer that can be passed to *ClockTime()*. If an error occurs, the function returns -1 and sets *errno*.

ClockId_r() An integer that can be passed to *ClockTime()*. This function does **NOT** set *errno*. If an error occurs, the function returns the negative of a value from the Errors section.

Errors:

ESRCH The *pid* and/or *tid* don't exist.

Examples:

Here's how you can determine how busy a system is:

```
id = ClockId(1, 1);
for(;;) {
    ClockTime(id, NULL, &start);
    sleep(1);
    ClockTime(id, NULL, &stop);
    printf("load = %f%%\n", (1000000000.0 - (stop-start)) / 1000000.0);
}
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ClockTime(), clock_getcpuclockid(), pthread_getcpuclockid()

Synopsis:

```
#include <sys/neutrino.h>

int ClockPeriod( clockid_t id,
                 const struct _clockperiod * new,
                 struct _clockperiod * old,
                 int reserved );

int ClockPeriod_r( clockid_t id,
                   const struct _clockperiod * new,
                   struct _clockperiod * old,
                   int reserved );
```

Arguments:

<i>id</i>	The clock ID of the clock. This must be CLOCK_REALTIME, which is the ID of the clock that maintains the system time.
<i>new</i>	NULL, or a pointer to a <code>_clockperiod</code> structure that contains the period to set the clock to. This structure contains at least the following members: <ul style="list-style-type: none">• <code>unsigned long nsec</code> — the period of the clock, in nanoseconds.• <code>long fract</code> — reserved for future fractional nanoseconds. Set this member to zero.
<i>old</i>	NULL, or a pointer to a <code>_clockperiod</code> structure where the function can store the current period (before being changed by a non-NULL <i>new</i>).
<i>reserved</i>	Set this argument to 0.

Library:

`libc`

Description:

You can use the *ClockPeriod()* and *ClockPeriod_r()* kernel calls to get or set the clock period of the clock.

These functions are identical except in the way they indicate errors. See the Returns section for details.



You need to have superuser privileges to set the clock period.

All the *timer_**(*)* calls operate with an accuracy no better than the clock period. Every moment within the Neutrino microkernel is referred to as a *tick*. A tick's initial length is determined by the clock rate of your processor:

CPU clock speed:	Default value:
≥ 40MHz	1 millisecond
< 40MHz	10 milliseconds

Since a very small ticksize imposes an interrupt load on the system, and can consume all available processor cycles, the kernel call limits how small a period can be specified. The lowest clock period that can currently be set on any machine is 10 microseconds.

If an attempt is made to set a value that the kernel believes to be unsafe, the call fails with an EINVAL. The timeslice rate (for “round-robin” and “other” scheduling policies) is always four times the clock period (this isn't changeable).

Blocking states

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

- ClockPeriod()* If an error occurs, this function returns -1 and sets *errno*. Any other value returned indicates success.
- ClockPeriod_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function can return any value in the Errors section.

Errors:

- EFAULT A fault occurred when the kernel tried to access the buffers provided.
- EINVAL Invalid clock ID. A period was set which wasn't in a range considered safe.
- EPERM The process tried to change the period of the clock without having superuser capabilities.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ClockAdjust()

ClockTime()*, *ClockTime_r()

© 2004, QNX Software Systems Ltd.

Get or set a clock

Synopsis:

```
#include <sys/neutrino.h>

int ClockTime( clockid_t id,
               const uint64_t * new,
               uint64_t * old );

int ClockTime_r( clockid_t id,
                 const uint64_t * new,
                 uint64_t * old );
```

Arguments:

id The clock ID. This must be CLOCK_REALTIME or CLOCK_MONOTONIC, which is the ID of the clock that maintains the system time, or the clock ID that's returned by *ClockId()*.

new NULL, or a pointer to the absolute time, in nanoseconds, to set the clock to.

old NULL, or a pointer to a location where the function can store the current time (before being changed by a non-NULL *new*).

Library:

libc

Description:

You can use these kernel calls to get or set the system clock specified by *id*. The clock ID, CLOCK_REALTIME or CLOCK_MONOTONIC, maintains the system time.

The *ClockTime()* and *ClockTime_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

If *new* isn't NULL, then it contains the absolute time, in nanoseconds, to set the system clock to. This affects the software clock maintained

by the system. It won't change any underlying hardware clock that maintains the time when the system's power is turned off.

Once set, the system time increments by some number of nanoseconds, based on the resolution of the system clock. You can query or change this resolution by using the *ClockPeriod()* kernel call.



You need to have superuser privileges to set the clock.

Blocking states

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

ClockTime() If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

ClockTime_r() EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function returns a value in the Errors section.

Errors:

EFAULT	A fault occurred when the kernel tried to access the buffers provided.
EINVAL	The clock ID isn't CLOCK_REALTIME or CLOCK_MONOTONIC.
EPERM	The process tried to change the time without having superuser capabilities.
ESRCH	The process associated with this request doesn't exist.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ClockAdjust(), *ClockPeriod()*

Synopsis:

```
#include <unistd.h>

int close( int filedes );
```

Arguments:

filedes The file descriptor of the file you want to close. This can be a file descriptor returned by a successful call to *accept()*, *creat()*, *dup()*, *dup2()*, *fcntl()*, *modem_open()*, *open()*, *shm_open()*, *socket()* or *sopen()*.

Library:

libc

Description:

The *close()* function closes the file specified by the given file descriptor.

Returns:

Zero for success, or -1 if an error occurs (*errno* is set).

Errors:

EBADF	Invalid file descriptor <i>filedes</i> .
EINTR	The <i>close()</i> call was interrupted by a signal.
EIO	An I/O error occurred while updating the directory information.
ENOSPC	A previous buffered write call has failed.
ENOSYS	The <i>close()</i> function isn't implemented for the filesystem specified by <i>filedes</i> .

Examples:

```
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main( void )
{
    int filedes;

    filedes = open( "file", O_RDONLY );
    if( filedes != -1 ) {
        /* process file */
        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

accept(), *creat()*, *dup()*, *dup2()*, *errno*, *fcntl()*, *modem_open()*, *open()*, *shm_open()*, *socket()*, *sopen()*

Synopsis:

```
#include <dirent.h>

int closedir( DIR * dirp );
```

Arguments:

dirp A directory pointer for the directory you want to close.

Library:

`libc`

Description:

The *closedir()* function closes the directory specified by *dirp*, and frees the memory allocated by *opendir()*.



The result of using a directory stream after calling one of the *exec*()* or *spawn*()* family of functions is undefined. After a call to the *fork()* function, either the parent or the child (but not both) may continue processing the directory stream using the *readdir()* and *rewinddir()* functions. If both the parent and child processes use these functions, the result is undefined. Either or both processes may call the *closedir()* function.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EBADF The *dirp* argument doesn't refer to an open directory stream.

EINTR The *closedir()* call was interrupted by a signal.

Examples:

Get a list of files contained in the directory `/home/kenny`:

```
#include <stdio.h>
#include <dirent.h>
#include <stdlib.h>

int main( void )
{
    DIR *dirp;
    struct dirent *direntp;

    dirp = opendir( "/home/kenny" );
    if( dirp != NULL ) {
        for(;;) {
            direntp = readdir( dirp );
            if( direntp == NULL ) {
                break;
            }

            printf( "%s\n", direntp->d_name );
        }

        closedir( dirp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

*errno, opendir(), readdir(), readdir_r(), rewinddir(), seekdir(),
telldir()*

closelog()

© 2004, QNX Software Systems Ltd.

Close the system log

Synopsis:

```
#include <syslog.h>

void closelog( void );
```

Library:

libc

Description:

The *closelog()* function closes the connection to **syslogd**.

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

openlog(), *setlogmask()*, *syslog()*, *vsyslog()*
logger, **syslogd** in the *Utilities Reference*

Synopsis:

```
#include <process.h>

int _cmdfd( void );
```

Library:

libc

Description:

This function returns a file descriptor for the executable file.

Returns:

A file descriptor for the executable file.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

_cmdname(), _progrname

_cmdname()

© 2004, QNX Software Systems Ltd.

Find the path used to invoke the current process

Synopsis:

```
#include <process.h>

char * _cmdname( char * buff );
```

Arguments:

buff A pointer to a buffer in which the function can store the path. To determine the size required for the buffer, call *fpathconf()* or *pathconf()* with an argument of `_PC_PATH_MAX`.

Library:

`libc`

Description:

The *_cmdname()* function determines the full path that the current process was invoked from, and stores it in the buffer specified by *buff*.

Returns:

0 Success.
-1 An error occurred.

Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <unistd.h>
#include <limits.h>
#include <process.h>

int main( void )
{
    size_t maximum_path;
    char *buff;

    maximum_path = (size_t) pathconf( "/", _PC_PATH_MAX );
    buff = (char* )malloc( maximum_path );
```

```
if( _cmdname( buff ) ) {
    printf( "I'm \"%s\".\n", buff );
} else {
    perror( "_cmdname() failed" );
    free (buff);
    return EXIT_FAILURE;
}

free (buff);
return EXIT_SUCCESS;
}
```

If this code is compiled into an executable named **foo**:

```
# ls -F /home/xyzyz/bin/foo
foo*
# /home/xyzyz/bin/foo
I'm "/home/xyzyz/bin/foo".
```

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

basename(), _cmdfd(), _progrname

confstr()

© 2004, QNX Software Systems Ltd.

Get configuration-defined string values

Synopsis:

```
#include <unistd.h>

size_t confstr( int name,
                char * buf,
                size_t len );
```

Arguments:

name The system variable to query; see below.

buf A pointer to a buffer in which the function can store the value of the system variable.

len The length of the buffer, in bytes.

Library:

`libc`

Description:

The *confstr()* functions lets applications get or set configuration-defined string values. This is similar to the *sysconf()* function, but you use it to get string values, rather than numeric values. By default, the function queries and returns values in the system.

The *name* argument represents the system variable to query. The values are defined in `<confname.h>`; at least the following *name* values are valid:

`_CS_ARCHITECTURE` The name of the instruction set architecture for this node's CPU(s).

`_CS_DOMAIN` The domain name.

`_CS_HOSTNAME` The name of this node in the network.



A hostname can consist only of letters, numbers, and hyphens, and must not start or end with a hyphen. For more information, see *RFC 952*.

`_CS_HW_PROVIDER`

The name of the hardware manufacturer.

`_CS_HW_SERIAL` Serial number associated with the hardware.

`_CS_LIBPATH` A value similar to the **LD_LIBRARY_PATH** environment variable that finds all standard libraries.

`_CS_MACHINE` This node's hardware type.

`_CS_PATH` A value similar to the **PATH** environment variable that finds all standard utilities.

`_CS_RELEASE` The current OS release level.

`_CS_RESOLVE` The contents of the `resolv.conf` file, excluding the domain name.

`_CS_SRPC_DOMAIN`

The secure RPC domain.

`_CS_SYSNAME` The operating system name.

`_CS_TIMEZONE` Time zone string (**TZ** style)

`_CS_VERSION` The current OS version number.

The configuration-defined value is returned in the buffer pointed to by *buf*, and will be $\leq len$ bytes long, including the terminating NULL.

To find out the length of a configuration-defined value, call *confstr()* with *buf* set to NULL and *len* set to 0.

To set a configuration value:

- OR your value to be defined (i.e. `_CS_HOSTNAME`) to `_CS_SET`
- put this value in a NULL-terminated string
- Set the value of *len* to 0

Returns:

A nonzero value (if a “get” is done, the value is the length of the configuration-defined value), or 0 if an error occurs (*errno* is set).

You can compare the *confstr()* return value against *len* to see if the configuration-defined value was truncated when retrieving a value, (this can't be done when setting a value).

Errors:

`EINVAL` The *name* argument isn't a valid configuration-defined value.

Examples:

Print information similar to that returned by the *uname()* function:

```
#include <unistd.h>
#include <stdio.h>
#include <limits.h>

#define BUFF_SIZE (256 + 1)

int main( void )
{
    char buff[BUFF_SIZE];

    if( confstr( _CS_SYSNAME, buff, BUFF_SIZE ) > 0 ) {
        printf( "System name: %s\n", buff );
    }

    if( confstr( _CS_HOSTNAME, buff, BUFF_SIZE ) > 0 ) {
        printf( "Host name: %s\n", buff );
    }

    if( confstr( _CS_RELEASE, buff, BUFF_SIZE ) > 0 ) {
        printf( "Release: %s\n", buff );
    }
}
```

```
if( confstr( _CS_VERSION, buff, BUFF_SIZE ) > 0 ) {
    printf( "Version: %s\n", buff );
}

if( confstr( _CS_MACHINE, buff, BUFF_SIZE ) > 0 ) {
    printf( "Machine: %s\n", buff );
}

if( confstr( _CS_SET | _CS_HOSTNAME, "myhostname", 0 ) != 0 ) {
    printf( "Hostname set to: %s\n", "myhostname" );
}

return 0;
}
```

Classification:

POSIX 1003.1a

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *confstr()* function is part of a draft standard; its interface and/or behavior may change in the future.

See also:

pathconf(), *sysconf()*

getconf, **setconf** in the *Utilities Reference*

“Configuration strings” in the Configuring Your Environment chapter
of the Neutrino *User’s Guide*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>

int connect( int s,
             const struct sockaddr * name,
             socklen_t namelen );
```

Arguments:

s The descriptor of the socket on which to initiate the connection.

name The name of the socket to connect to for a SOCK_STREAM connection.

namelen The length of the *name*, in bytes.

Library:

libsocket

Description:

The *connect()* function establishes the connection according to the socket type specified by *s*:

SOCK_DGRAM

Specifies the peer that the socket is to be associated with. This address is the one that datagrams are to be sent to, and the only one that datagrams are to be received from.

SOCK_STREAM

This call attempts to make a connection to another socket. The other socket is specified by *name*, which is an address in the communications space of that socket. Each communications space interprets *name* in its own way.

Stream sockets may successfully connect only once, whereas datagram sockets may use *connect()* multiple times to change their association. Datagram sockets may dissolve the association by connecting to an invalid address, such as a null address.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EADDRINUSE The address is already in use.
- EADDRNOTAVAIL
The specified address isn't available on this machine.
- EAFNOSUPPORT
Addresses in the specified address family cannot be used with this socket.
- EALREADY The socket is nonblocking; a previous connection attempt hasn't yet been completed.
- EBADF Invalid descriptor *s*.
- ECONNABORTED
The *connect()* was terminated under software control.
- ECONNREFUSED
The attempt to connect was forcefully rejected.
- EFAULT The *name* parameter specifies an area outside the process address space.
- EHOSTUNREACH
No route to host; the host system can't be reached.

EINPROGRESS	The socket is nonblocking; the connection can't be completed immediately. It's possible to do a <i>select()</i> for completion by selecting the socket for writing.
EISCONN	The socket is already connected.
ENETUNREACH	The network isn't reachable from this host.
ETIMEDOUT	The attempt to establish a connection timed out; no connection was made.



Protocols such as TCP do not allow further connection requests on a socket after an ECONNREFUSED error. In such a situation, the socket must be closed and a new one created before a subsequent attempt for connection is made.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

ICMP, IP, TCP, and UDP protocols

accept(), *bind()*, *getsockname()*, *nbaconnect()*, *select()*, *socket()*

ConnectAttach(), ConnectAttach_r() © 2004, QNX Software Systems

Ltd.

Establish a connection between a process and a channel

Synopsis:

```
#include <sys/neutrino.h>

int ConnectAttach( uint32_t nd,
                  pid_t pid,
                  int chid,
                  unsigned index,
                  int flags );

int ConnectAttach_r( uint32_t nd,
                    pid_t pid,
                    int chid,
                    unsigned index,
                    int flags );
```

Arguments:

- nd* The node descriptor of the node on which the process that owns the channel is running; see “Node descriptors,” below.
- pid* The process ID of the owner of the channel. If *pid* is zero, the calling process is assumed.
- chid* The channel ID, returned by *ChannelCreate()*, of the channel to connect to the process.
- index* The lowest acceptable connection ID.



Treating a connection as a file descriptor can lead to unexpected behavior. Therefore, you should OR `_NTO_SIDE_CHANNEL` into *index* when you create a connection. If you do this, the connection ID is returned from a different space than file descriptors; the ID is greater than any valid file descriptor.

Once created there's no difference in the use of the messaging primitives on this ID. The C library creates connections at various times without `_NTO_SIDE_CHANNEL` (e.g. during *open()*), however, it's unlikely that any applications would want to call it this way.

flags If *flags* contains `_NTO_COF_CLOEXEC`, the connection is closed when your process calls an *exec*()* function to start a new process.

Library:

`libc`

Description:

The *ConnectAttach()* and *ConnectAttach_r()* kernel calls establish a connection between the calling process and the channel specified by *chid* owned by the process specified by *pid* on the node specified by *nd*. Any function that passes a node descriptor can use either the value 0 or the constant `ND_LOCAL_NODE` to refer to the local node.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The return value is a connection ID, which is a small `int` representing the connection. The system returns the first available connection ID starting at the value specified by the *index* argument. Any thread in the calling process can use either *MsgSendv()* to send messages or *MsgSendPulse()* to send pulses over the connection. The connection ID is used directly as a POSIX file descriptor (*fd*) when communicating with I/O Resource managers such as a filesystem manager.

If you don't OR `_NTO_SIDE_CHANNEL` into *index*, this behavior might result:

- If file descriptor 0 is in use, file descriptor 1 isn't in use, and you call `ConnectAttach()` with 0 specified for *index*, a connection ID of 1 is returned.

File descriptor 1 (i.e. connection ID 1) is used as *stdout*, which is what `printf()` writes to. If your process makes any calls to `printf()`, NULL-terminated character strings are sent to the channel that you've connected to. Similar situations can happen with connection IDs 0 (*stdin*) and 2 (*stderr*).

- Depending on how a child process is created, it may inherit the parent's file descriptors.

Since connections are treated like file descriptors, a connection created by the parent without `_NTO_SIDE_CHANNEL` in *index* and without `_NTO_COF_CLOEXEC` in *flags*, causes a child process to inherit that connection during process creation. This inheritance is done during process creation by duplicating file descriptors.

During duplication, an `_IO_DUP` message (with `0x115`) as the first 2 bytes) is sent to the receiver on the other side of the connection. The receiver won't be expecting this message.

If *index* has `_NTO_SIDE_CHANNEL` set, the *index* is ignored and the connection ID returned is the first available index in the `_NTO_SIDE_CHANNEL` space.

If a process creates multiple connections to the same channel, the system maintains a link count and shares internal kernel object resources for efficiency.

Connections are owned by the process and may be used simultaneously by any thread in the process. You can detach a connection by calling `ConnectDetach()`. If any threads are blocked on the channel (via `MsgSendv()`) at the time the connection is detached, the send fails and returns with an error.



Connections and connection IDs persist until you call *ConnectDetach()*, even if the other process dies.

The connection is strictly local (i.e. it doesn't resolve across the network) and is resolved on the first use of the connection ID.

Blocking states

These calls don't block.

Node descriptors

The *nd* (node descriptor) is a temporary numeric description of a remote node. For more information, see the Qnet Networking chapter of the *System Architecture* guide.

To:	Use this function:
Compare two <i>nd</i> objects	<i>ND_NODE_CMP()</i>
Convert a <i>nd</i> to text	<i>netmgr_ndtostr()</i>
Convert text to a <i>nd</i>	<i>netmgr_strtond()</i>

Returns:

The only difference between these functions is the way they indicate errors:

ConnectAttach()

A connection ID that's used by the message primitives. If an error occurs, the function returns -1 and sets *errno*.

ConnectAttach_r()

A connection ID that's used by the message primitives. This function does **NOT** set *errno*. If an error occurs, the function returns the negative of a value from the Errors section.

Errors:

EAGAIN	All kernel connection objects are in use.
ESRCH	The node indicated by <i>nd</i> , the process indicated by <i>pid</i> , or the channel indicated by <i>chid</i> doesn't exist.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ChannelCreate(), *ConnectDetach()*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *MsgSendPulse()*, *MsgSendv()*, *netmgr_remote_nd()*

ConnectClientInfo(), ConnectClientInfo_r()

Store information about a client connection

Synopsis:

```
#include <sys/neutrino.h>

int ConnectClientInfo( int scoid,
                      struct _client_info * info
                      int ngroups );

int ConnectClientInfo_r( int scoid,
                        struct _client_info * info
                        int ngroups );
```

Arguments:

- scoid* A server connection ID that identifies the client process that you want to get information about. This client is typically a process that's made a connection to the server to try to access a resource. You can get it from the `_msg_info` argument to *MsgReceivev()* or *MsgInfo()*.
- info* A pointer to a `_client_info` structure that the function can fill with information about the client. For more information, see below.
- ngroups* The size of the caller's *group*list in the credential part of the `_client_info` structure. If you make it smaller than `NGROUPS_MAX`, you might get information only about a subset of the groups.

Library:

`libc`

Description:

These calls get information about a client connection identified by *scoid*, and store it in the buffer that *info* points to.

The *ConnectClientInfo()* and *ConnectClientInfo_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

A server uses these functions to determine whether or not a client has permission to access a resource. For example, in a resource manager, it would be called on an *open()* connection request.

_client_info structure

The **_client_info** structure has at least the following members:

uint32_t *nd* The client's node ID.

pid_t *pid* The client's process ID.

struct _cred_info *cred*
The user and group ID credentials; see below.

uint32_t nd

The *nd* (node descriptor) is a temporary numeric description of a remote node. For more information, see the Qnet Networking chapter of the *System Architecture* guide.

To:	Use this function:
Compare two <i>nd</i> objects	<i>ND_NODE_CMP()</i>
Convert a <i>nd</i> to text	<i>netmgr_ndtostr()</i>
Convert text to a <i>nd</i>	<i>netmgr_strtond()</i>

_cred_info structure

The *cred* member of the **_client_info** is a **_cred_info** structure that includes at least the following members:

uid_t *ruid* The real user ID of the sending process.

uid_t *euid* The effective user ID of the sending process.

uid_t *suid* The saved user ID of the sending process.

gid_t *rgid* The real group ID of the sending process.

gid_t *egid* The effective group ID of the sending process.

gid_t *sgid* The saved group ID of the sending process.

uint32_t *ngroups*
 The number of groups actually stored in *grouplist*.

gid_t *grouplist*[NGROUPS_MAX]
 The supplementary group IDs of the sending process.

The *ngroups* argument to *ConnectClientInfo()* indicates the size of the *grouplist* array. If the group array size is zero, the *ngroups* member of the *_cred_info* is set to the number of groups available.

Returns:

The only difference between these functions is the way they indicate errors:

ConnectClientInfo()

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

ConnectClientInfo_r()

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function can return any value in the Errors section.

Errors:

EFAULT A fault occurred when the kernel tried to access the buffers provided.

EINVAL Process doesn't have a connection *scoid*.

ConnectClientInfo(), ConnectClientInfo_r()

© 2004, QNX

Software Systems Ltd.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*ConnectServerInfo(), _msg_info, MsgInfo(), MsgReceive(),
ND_NODE_CMP(), netmgr_ndtostr(), netmgr_remote_nd(),
netmgr_strtond()*

ConnectDetach()*, *ConnectDetach_r()

Break a connection between a process and a channel

Synopsis:

```
#include <sys/neutrino.h>

int ConnectDetach( int coid );

int ConnectDetach_r( int coid );
```

Arguments:

coid The connection ID of the connection you want to break.

Library:

`libc`

Description:

The *ConnectDetach()* and *ConnectDetach_r()* kernel calls detach the connection specified by the *coid* argument. If any threads are blocked on the connection (*MsgSendv()*) at the time the connection is detached, the send fails and returns with an error.

These functions are identical except in the way they indicate errors. See the Returns section for details.

Blocking states

These calls don't block.

Returns:

The only difference between these functions is the way they indicate errors:

ConnectDetach()

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

ConnectDetach(), *ConnectDetach_r()*

© 2004, QNX Software

Systems Ltd.

ConnectDetach_r()

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, the function returns a value in the Errors section.

Errors:

EINVAL The connection specified by *coid* doesn't exist.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ConnectAttach(), *MsgSendv()*

ConnectFlags()*, *ConnectFlags_r()

Modify the flags associated with a connection

Synopsis:

```
#include <sys/neutrino.h>

int ConnectFlags( pid_t pid,
                 int coid,
                 unsigned mask,
                 unsigned bits );

int ConnectFlags_r( pid_t pid,
                  int coid,
                  unsigned mask,
                  unsigned bits );
```

Arguments:

- pid* The ID of the process that the connection ID belongs to, or 0 for the current process.
- coid* The ID of the connection whose flags you want to modify.
- mask* A bitmap that indicates which bits are to be modified in the flags.
- bits* The new value of the flags. The flags currently defined include:
- `_NTO_COF_CLOEXEC` — close the connection if the process calls an *exec*()* function to start a new process.

Library:

`libc`

Description:

The *ConnectFlags()* and *ConnectFlags_r()* kernel calls modify flags associated with the specified connection. These kernel calls don't block.

These functions are identical except in the way they indicate errors. See the Returns section for details.

You need to initialize the bits that correspond to the flag in both the *mask* and *bits* arguments:

- If the bit in the *mask* is 1, and the bit in the *bits* is 1, the function turns the flag on.
- If the bit in the *mask* is 1, and the bit in the *bits* is 0, the function turns the flag off.
- If bit in the mask is 0, the function doesn't change the current value of the flag.

Returns:

The only difference between these functions is the way they indicate errors:

ConnectFlags()

The previous value of the flags associated with the connection. If an error occurs, the function returns -1 and sets *errno*.

ConnectFlags_r()

The previous value of the flags associated with the connection. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

Errors:

ESRCH The process ID is invalid or the connection ID can't be found.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ConnectAttach(), *fcntl()*

ConnectServerInfo(), ConnectServerInfo_r() © 2004, QNX

Software Systems Ltd.

Get information about a server connection

Synopsis:

```
#include <sys/neutrino.h>

int ConnectServerInfo( pid_t pid,
                      int coid,
                      struct _server_info* info );

int ConnectServerInfo_r( pid_t pid,
                        int coid,
                        struct _server_info* info );
```

Arguments:

pid The process ID of the owner of the connection.

coid The connection ID of the connection.

info A pointer to a `_server_info` structure where the function can store information about the connection. For more information, see below.

Library:

`libc`

Description:

The `ConnectServerInfo()` and `ConnectServerInfo_r()` kernel calls get information about the connection *coid* owned by process *pid*, and store it in the structure pointed to by *info*. If the process doesn't have a connection *coid*, the call scans for the next higher connection and returns it if present. Otherwise, -1 is returned. If you wish to check for the existence of an exact connection, you must compare the returned connection with the *coid* you requested.

These functions are identical except in the way they indicate errors. See the Returns section for details.

_server_info structure

The `_server_info` structure that *info* points to includes at least the following members:

<code>uint32_t nd</code>	The server's node ID.
<code>pid_t pid</code>	The server's process ID.
<code>int32_t chid</code>	The server's channel ID.
<code>int32_t scoid</code>	The server's connection ID.

uint32_t nd

The *nd* (node descriptor) is a temporary numeric description of a remote node. For more information, see the Qnet Networking chapter of the *System Architecture* guide.

To:	Use this function:
Compare two <i>nd</i> objects	<code>ND_NODE_CMP()</code>
Convert a <i>nd</i> to text	<code>netmgr_ndtostr()</code>
Convert text to a <i>nd</i>	<code>netmgr_strtond()</code>

Returns:

The only difference between these functions is the way they indicate errors:

ConnectServerInfo()

A matched *coid*. If an error occurs, the function returns -1 and sets *errno*.

ConnectServerInfo(), ConnectServerInfo_r()

© 2004, QNX

Software Systems Ltd.

ConnectServerInfo_r()

A matched *coid*. This function does **NOT** set *errno*. If an error occurs, the function returns the negative of a value from the Errors section.

Errors:

EFAULT	A fault occurred when the kernel tried to access the buffers provided.
EINVAL	Process <i>pid</i> doesn't have a connection \geq <i>coid</i> .
ESRCH	The process indicated by <i>pid</i> doesn't exist.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ConnectAttach(), *ConnectClientInfo()*, *MsgInfo()*, *MsgReceivev()*, *ND_NODE_CMP()*, *netmgr_ndtostr()*, *netmgr_remote_nd()*, *netmgr_strtond()*

Synopsis:

```
#include <math.h>

double copysign ( double x,
                 double y );

float copysignf ( float x,
                 float y );
```

Arguments:

- x* The number to use the magnitude of.
- y* The number to use the sign of.

Library:

`libm`

Description:

The *copysign()* and *copysignf()* functions return the magnitude of *x* and the sign bit of *y*.

If *x* is NAN, the function produces NAN with the sign of *y*.

Returns:

The magnitude of *x* and the sign bit of *y*.

Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
    double a, b, c;

    a = 27.0;
```

```
b = -5;
c = copysign(a, b);
printf("The magnitude of %f and sign of %f gives %f\n",
      a, b, c);

return(0);
}
```

produces the output:

```
The magnitude of 27.000000 and sign of -5.000000 gives -27.000000
```

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

significand()

Synopsis:

```
#include <math.h>

double cos( double x );

float cosf( float x );
```

Arguments:

x The angle, in radians, for which you want to compute the cosine.

Library:

libm

Description:

These functions compute the cosine of *x* (specified in radians).



An argument with a large magnitude may yield results with little or no significance.

Returns:

The cosine of *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <math.h>
#include <stdio.h>
#include <stdlib.h>
```

```
int main( void )
{
    double value;

    value = cos( M_PI );
    printf( "value = %f\n", value );

    return EXIT_SUCCESS;
}
```

produces the output:

```
value = -1.000000
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

acos(), errno, sin(), tan()

Synopsis:

```
#include <math.h>

double cosh( double x );

float coshf( float x );
```

Arguments:

- x* The angle, in radians, for which you want to compute the hyperbolic cosine.

Library:

libm

Description:

These functions compute the hyperbolic cosine (specified in radians) of *x*. A range error occurs if the magnitude of *x* is too large.

Returns:

The hyperbolic cosine of *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", cosh(.5) );
}
```

```
    return EXIT_SUCCESS;  
}
```

produces the output:

```
1.127626
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *sinh()*, *tanh()*

Synopsis:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

int creat( const char* path,
           mode_t mode );

int creat64( const char* path,
             mode_t mode );
```

Arguments:

path The path of the file you want to open.

mode The access permissions that you want to use. For more information, see “Access permissions” in the documentation for *stat()*.

Library:

`libc`

Description:

The *creat()* and *creat64()* functions create and open the file specified by *path* with the given *mode*.

Calling *creat()* is the same as:

```
open( path, O_WRONLY | O_CREAT | O_TRUNC, mode );
```

Similarly, calling *creat64()* is the same as:

```
open64( path, O_WRONLY | O_CREAT | O_TRUNC | O_LARGEFILE, mode );
```

If *path* exists and is writable, it’s truncated to contain no data, and the existing *mode* setting isn’t changed.

If *path* doesn’t exist, it’s created with the access permissions specified by the *mode* argument. The access permissions for the file or directory are specified as a combination of the bits defined in `<sys/stat.h>`.

Returns:

A file descriptor on success, or -1 if an error occurs (*errno* is set).

Errors:

EACCES	Indicates one of the following permission problems: <ul style="list-style-type: none">• Search permission is denied for one of the components in the <i>path</i>.• The file specified by <i>path</i> exists, and the permissions specified by <i>mode</i> are denied.• The file specified by <i>path</i> doesn't exist, and the file couldn't be created because write permission is denied for the parent directory.
EBADFSYS	While attempting to open <i>path</i> , the file itself or a component of its path prefix was found to be corrupted. A system failure — from which no automatic recovery is possible — occurred while the file was being written to or while the directory was being updated. The filesystem must be repaired before proceeding.
EBUSY	The file specified by <i>path</i> is a block special device that's already open for writing, or <i>path</i> names a file on a filesystem mounted on a block special device that is already open for writing.
EINTR	The call to <i>creat()</i> was interrupted by a signal.
EISDIR	The file specified by <i>path</i> is a directory and the file creation flags specify write-only or read/write access.
ELOOP	Too many levels of symbolic links.
EMFILE	This process is using too many file descriptors.
ENAMETOOLONG	The length of <i>path</i> exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> .

ENFILE	Too many files are currently open in the system.
ENOENT	Either the path prefix doesn't exist, or the <i>path</i> argument points to an empty string.
ENOSPC	The directory or filesystem that would contain the new file doesn't have enough space available to create a new file.
ENOSYS	The <i>creat()</i> function isn't implemented for the filesystem specified by <i>path</i> .
ENOTDIR	A component of the path prefix isn't a directory.
EROFS	The file specified by <i>path</i> resides on a read-only filesystem.

Examples:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdlib.h>

int main( void )
{
    int filedes;

    filedes = creat( "file",
                   S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP );
    if( filedes != -1 ) {
        /* process file */

        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

creat() is POSIX 1003.1; *creat64()* is for large-file support

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

chsize(), *close()*, *dup()*, *dup2()*, *eof()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *fcntl()*, *fileno()*, *fstat()*, *isatty()*, *lseek()*, *open()*, *read()*, *sopen()*, *stat()*, *tell()*, *write()*, *umask()*

Synopsis:

```
#include <unistd.h>

char * crypt( const char * key,
              const char * salt );
```

Arguments:

- key* A NUL-terminated string (normally a password typed by a user).
- salt* A two-character string chosen from the set [a-zA-Z0-9./]. This function doesn't validate the values for *salt*, and values outside this range may cause undefined behavior. This string is used to perturb the algorithm in one of 4096 different ways.

Library:

libc

Description:

The *crypt()* function performs password encryption. It's based on the Data Encryption Standard algorithm, and also includes code to deter key search attempts.



This function checks only the first eight characters of *key*.

You can obtain a 56-bit key by taking the lowest 7 bits of *key*. The 56-bit key is used to repeatedly encrypt a constant string (usually all zeroes).

Returns:

A pointer to the 13-character encrypted value, or NULL on failure. The first two characters of the encrypted value are the *salt* itself.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

The return value points to static data that's overwritten by each call to *crypt()*.

See also:

encrypt(), *getpass()*, *qnx_crypt()*, *setkey()*

login in the *Utilities Reference*

Copyright © MINIX Operating System

Synopsis:

```
#include <stdio.h>

char * ctermid( char * s );
```

Arguments:

s NULL, or a pointer to a buffer in which the function can store the path name of the controlling terminal. This string should be at least `L_ctermid` characters long (see `<stdio.h>`).

Library:

`libc`

Description:

The `ctermid()` function generates a string that contains the path name of the current controlling terminal for the calling process.



If the argument *s* is NULL, the string is built in a static buffer, and the function returns a pointer to the buffer.

Returns:

A pointer to the path name of the controlling terminal, or a pointer to a null string if the function can't locate the controlling terminal.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    printf( "Controlling terminal is %s\n", ctermid( NULL ) );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Read the <i>Caveats</i>

Caveats:

The *ctermid()* function isn't thread-safe if the *s* argument is NULL.

See also:

setsid(), *ttyname()*

Synopsis:

```
#include <time.h>

char* ctime( const time_t* timer );

char* ctime_r( const time_t* timer,
               char* buf );
```

Arguments:

timer A pointer to a `time_t` object that contains the time that you want to convert to a string.

buf (*ctime_r()* only) A buffer in which *ctime_r()* can store the resulting string. This buffer must be large enough to hold at least 26 characters.

Library:

`libc`

Description:

The *ctime()* and *ctime_r()* functions convert the time pointed to by *timer* to local time and formats it as a string containing exactly 26 characters in the form:

```
Tue May 7 10:40:27 2002\n\0
```

This function: Is equivalent to calling:

<i>ctime()</i>	<code>asctime(localtime (timer));</code>
<i>ctime_r()</i>	<code>asctime_r(localtime (timer), buf)</code>

The *ctime()* function places the result string in a static buffer that's reused each time *ctime()* or *asctime()* is called. The result string for *ctime_r()* is contained in the buffer pointed to by *buf*.

All fields have a constant width. The newline character '`\n`' and NUL character '`\0`' occupy the last two positions of the string.

Whenever the *ctime()* or *ctime_r()* functions are called, the *tzset()* function is also called.

The calendar time is usually obtained by using the *time()* function. That time is Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time (GMT)).

You typically set the time on the computer with the `date` command to reflect Coordinated Universal Time (UTC), and then use the `TZ` environment variable or `_CS_TIMEZONE` configuration string to establish the local time zone. For more information, see “Setting the time zone” in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

Returns:

A pointer to the string containing the formatted local time, or NULL if an error occurs.

Classification:

ctime() is ANSI, POSIX 1003.1; *ctime_r()* is POSIX 1003.1

ctime()

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	No

ctime_r()**Safety**

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *asctime()* and *ctime()* functions place their results in a static buffer that's reused for each call to *asctime()* or *ctime()*.

See also:

asctime(), *asctime_r()*, *clock()*, *difftime()*, *gmtime()*, *localtime()*, *localtime_r()*, *mktime()*, *strftime()*, *time()*, *tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User's Guide*

daemon()

© 2004, QNX Software Systems Ltd.

Run a process in the background

Synopsis:

```
#include <stdlib.h>

int daemon( int nochdir,
            int noclose );
```

Arguments:

nochdir If this argument is 0 — the current working directory to the root (/) is changed.

noclose If this argument is 0 — standard input, standard output and standard error to `/dev/null` are redirected.

Library:

`libc`

Description:

The *daemon()* function allows programs to detach themselves from the controlling terminal and run in the background as system daemons.

This function calls *fork()* and *setsid()*.



The controlling terminal behaves like the behavior intended in Unix System Version, Release 4. An *open()* on a terminal device not already associated with another session will cause the device to become the controlling terminal for that process.

Returns:

Zero for success, or -1 if an error occurs (*errno* is set).

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	No

Caveats:

Currently, *daemon()* is supported only in single-threaded applications. If you create a thread and then call *daemon()*, the function returns -1 and sets *errno* to ENOSYS.

See also:

fork(), *procmgr_daemon()*, *setsid()*

daylight

© 2004, QNX Software Systems Ltd.

Indicator of support for daylight saving time in the locale

Synopsis:

```
#include <time.h>

unsigned int daylight;
```

Description:

This global variable has a value of 1 when daylight saving time is supported in this locale, and 0 otherwise. Whenever you call a time function, *tzset()* is called to set the variable, based on the current time zone.

Classification:

QNX Neutrino

See also:

timezone, tzname, tzset()

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*

Synopsis:

```
#include <sys/neutrino.h>

void DebugBreak( void );
```

Library:

libc

Description:

The *DebugBreak()* kernel call activates the process debugger if you're debugging the calling process. If not, it sends a SIGTRAP signal to the process.

Blocking states

None.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

If you call *DebugBreak()* from an interrupt handler, it'll activate the *kernel* debugger (if it's present in your boot image) or send the process a SIGTRAP signal.

See also:

DebugKDBreak(), DebugKDOutput()

Synopsis:

```
#include <sys/neutrino.h>

void DebugKDBreak( void );
```

Library:

libc

Description:

The *DebugKDBreak()* kernel call activates the kernel debugger if it's present in your boot image. If not, nothing happens.

Blocking states

None.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

DebugBreak(), *DebugKDOutput()*

DebugKDOutput()

© 2004, QNX Software Systems Ltd.

Print text with the kernel debugger

Synopsis:

```
#include <sys/neutrino.h>

void DebugKDOutput( const char* str,
                   size_t size );
```

Arguments:

str The string that you want to print.

size The number of characters to print.

Library:

libc

Description:

The *DebugKDBreak()* kernel call causes the kernel debugger to print *size* characters from *str* if the kernel debugger is present in your boot image. If it isn't in your boot image, nothing happens.

When, where, and how the kernel debugger displays this message depends on which host debugger you're using.

Blocking states

None.

Classification:

QNX Neutrino

Safety

Cancellation point No

Interrupt handler No

continued...

Safety

Signal handler	Yes
Thread	Yes

See also:

DebugBreak(), DebugKDBreak()

delay()

© 2004, QNX Software Systems Ltd.

Suspends a calling thread for a given length of time

Synopsis:

```
#include <unistd.h>

unsigned int delay( unsigned int duration );
```

Arguments:

duration The number of milliseconds for which to suspend the calling thread from execution.

Library:

libc

Description:

The *delay()* function suspends the calling thread for *duration* milliseconds.



The suspension time may be greater than the requested amount, due to the scheduling of other, higher-priority threads by the system.

Returns:

0 for success, or the number of unslept milliseconds if interrupted by a signal.

Errors:

If an error occurs, *errno* is set to:

EAGAIN No timer resources were available to satisfy the request.

Examples:

```
#include <unistd.h>
#include <stdlib.h>

void play_sound( void )
{
    ...
}

void stop_sound( void )
{
    ...
}

int main( void )
{
    play_sound();
    delay( 500 ); /* delay for 1/2 second */
    stop_sound();

    return EXIT_SUCCESS;
}
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:*alarm(), errno, nanosleep(), nap(), napms(), sleep()*

devctl()

Control a device

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>
#include <devctl.h>

int devctl( int filedes,
            int dcmd,
            void * dev_data_ptr,
            size_t n_bytes,
            int * dev_info_ptr );
```

Arguments:

<i>filedes</i>	A file descriptor that you obtained by opening the device.
<i>dcmd</i>	<p>A device-specific command for the process managing the open device. The set of valid device-control commands, the associated data interpretation, the returned <i>dev_info_ptr</i> values, and the effect of the command on the device all depend on the device driver.</p> <p>For specific commands, see the <code><sys/dcmd_*.h></code> header files; for general information, see “Device-control commands,” below.</p>
<i>dev_data_ptr</i>	<p>Depending on the command, this argument is one of:</p> <ul style="list-style-type: none">• a pointer to a buffer containing data to be passed to the driver• a receiving area for data coming from the driver• both of the above• NULL.
<i>n_bytes</i>	The size of the data to be sent to the driver, or the maximum size of the data to be received from the driver. <i>MsgSend()</i> is used to transfer the data.

dev_info_ptr A pointer to a location that the device to can use to return additional status information instead of just success or failure. The data returned via *dev_info_ptr* depends on the device driver.

Library:

`libc`

Description:

The *devctl()* function sends the device-specific command *dcmd* to the process managing the device opened as *filedes*. For example, you can send commands to specify properties for devices such as keyboards, sound cards or serial ports.

Device-control commands

Use these macros to set up the device-control commands:

_DIOF(class, cmd, data)

Get information from the device.

_DION(class, cmd)

A command with no associated data.

_DIOT(class, cmd, data)

Pass information to the device.

_DIOTF(class, cmd, data)

Pass some information to the device, and get some from it.

The arguments to these macros are:

class The major category for the command. The device-control commands are divided into the following classes to make organization easier:

- `_DCMD_ALL` — Common (all I/O servers).

- `_DCMD_CAM` — Low-level (Common Access Method) devices, such as disks or CD-ROMs.
- `_DCMD_CHR` — Character devices.
- `_DCMD_FSYS`, `_DCMD_BLK` — Filesystem/block I/O managers.
- `_DCMD_INPUT` — Input devices.
- `_DCMD_IP` — Internet Protocol.
- `_DCMD_MEM` — Memory card.
- `_DCMD_MISC` — Miscellaneous commands.
- `_DCMD_MIXER` — Mixer (Audio).
- `_DCMD_NET` — Network devices.
- `_DCMD_PHOTON` — Photon.
- `_DCMD_PROC` — Process manager.

cmd The specific command in the class.

data The type of data to pass to and/or from the device. The *dev_data_ptr* argument to *devctl()* must be a pointer to this type of data, and *n_bytes* is usually the size of this type of data.



The size of the structure that's passed as the last field to the `_DIO*` macros **must** be less than $2^{14} == 16\text{K}$. Anything larger than this interferes with the upper two directional bits.

Resource managers can use the following macros, which are defined in `<devctl.h>`, when handling commands:

get_device_command(cmd)

Extract the class and the specific device command from *cmd* (i.e. strip off the data type and the direction).

get_device_direction(cmd)

Get the direction of the command (DEVDIR_TO, DEVDIR_FROM, DEVDIR_TOFROM, or DEVDIR_NONE).

Returns:

EOK	Success.
EAGAIN	The <i>devctl()</i> command couldn't be completed because the device driver was in use by another process, or the driver was unable to carry out the request due to an outstanding command in progress.
EBADF	Invalid open file descriptor, <i>filedes</i> .
EINTR	The <i>devctl()</i> function was interrupted by a signal.
EINVAL	The device driver detected an error in <i>dev_data_ptr</i> or <i>n_bytes</i> .
EIO	The <i>devctl()</i> function couldn't complete because of a hardware error.
ENOSYS	The device doesn't support the <i>dcmd</i> command.
ENOTTY	The <i>dcmd</i> argument isn't a valid command for this device.
EPERM	The process doesn't have sufficient permission to carry out the requested command.

Examples:

Example 1: Setting RTS on a serial port

Here's a quick example of setting and unsetting RTS (Request to Send) on a serial port:

```
/* For "devctl()" */
#include <devctl.h>
#include <sys/dcmd_chr.h>

/* For "open()" */
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

/* For Errors */
#include <stdlib.h>
#include <stdio.h>
```

```
int check_RTS(int fd);

int main(void)
{
    int data = 0, fd, error;

    if((fd = open ("/dev/ser2", O_RDONLY)) == -1)
    {
        fprintf(stderr, "Error with open() on /dev/ser2. Make sure exists.\n");
        perror (NULL);
        exit(EXIT_FAILURE);
    }

    check_RTS(fd);

    /* Let's turn ON RTS now. */
    data = _CTL_RTS_CHG | _CTL_RTS;

    if (error = devctl (fd, DCMD_CHR_SERCTL, &data, sizeof(data), NULL))
    {
        fprintf(stderr, "Error setting RTS: %s\n",
            strerror ( error ));
        exit(EXIT_FAILURE);
    }
    /* RTS should now be ON. */

    check_RTS(fd);

    sleep (2);

    /* Now let's turn RTS OFF. */
    data = _CTL_RTS_CHG | 0;

    if (error = devctl (fd, DCMD_CHR_SERCTL, &data, sizeof(data), NULL))
    {
        fprintf(stderr, "Error setting RTS: %s\n",
            strerror ( error ));
        exit(EXIT_FAILURE);
    }
    /* RTS should now be OFF. */

    check_RTS(fd);

    close(fd);

    return (1);
}

int check_RTS(int fd)
{
    int data = 0, error;

    /*
     * Let's see if RTS is set, tell devctl() we're requesting
     * line status information and devctl() then assigns data
     * the line status information for us. Too easy.
     */
}
```

```

if (error = devctl (fd, DCMD_CHR_LINESTATUS, &data,
                  sizeof(data), NULL))
{
    fprintf(stderr, "Error setting RTS: %s\n",
           strerror ( error ));
    exit(EXIT_FAILURE);
}

if (data & _LINESTATUS_SER_RTS)

    printf("RTS is SET!\n");

else

    printf("RTS is NOT set\n");

return(1);
}

```

The two main areas of interest are the setting of *data* and the *devctl()* call. The data variable is used for both sending and receiving data.

When setting RTS, *data* is assigned a value that's sent to the device via *devctl()*.

If <i>data</i> equals:	RTS is turned:
<code>_CTL_RTS_CHG _CTL_RTS</code>	ON
<code>_CTL_RTS_CHG</code>	OFF

When checking to see if RTS is set, we call *devctl()* with *dcmd* set to the `DCMD_CHR_LINESTATUS` macro and *data* containing any value (zero is clean). The *devctl()* function returns with *data* containing the Line Status value. This then can be used to determine what lines are set on that device. In our example, we check against `_LINESTATUS_SER_RTS`.

To find out what values to use with different `DCMD_*` commands, look in the appropriate `<sys/dcmd_*.h>` header file. For example, you'll find macros for the following values under `DCMD_CHR_LINESTATUS` in `<sys/dcmd_chr.h>`:

- Serial Port (DTR, RTS, CTS, DSR, RI, CD)

- Keyboard (Scroll/Caps/Num Lock, Shift, CTRL, ALT)
- Parallel Port (No Error, Selected, Paper Out, No Tack, Not Busy)

The value that's in the header is a "bitwise &" with the value in *data* to see if the value is high for that line.

Example 2: Cycling through Caps Lock, Num Lock, and Scroll Lock

In the following example, we open the device `/dev/kbd` and we start applying changes to the Caps Lock, Scroll Lock, and Num Lock properties.

The key lines in this example are the same as in the last example; they focus around the data variable. This value is just a simple integer value that's passed into the `devctl()` function. The data variable is assigned its values by simply performing a bitwise OR to the predefined values in the `<usr/include/sys/dcmd_chr.h>` header. Note the values used in the bitwise OR:

- `_CONCTL_NUM_CHG` (Console Control Num Lock Change) ORed together with `_CONCTL_NUM` (Console Control Num Lock) turns on Num Lock.
- `_CONCTL_NUM_CHG` on its own turns off Num Lock.

If <i>data</i> equals:	Num Lock is turned:
<code>_CONCTL_NUM_CHG _CONCTL_NUM</code>	ON
<code>_CONCTL_NUM_CHG</code>	OFF

This also applies for the other either/or values in the `<dcmd_chr.h>` header.

```

/* For "devctl()" */
#include <devctl.h>
#include <sys/dcmd_chr.h>

/* For "open()" */
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

```

```

/* For Errors */
#include <stdlib.h>
#include <stdio.h>

int main(void)
{
    int data, fd, toggle = 1, error;

    /* Open the device we wish to manipulate. */
    if((fd = open ("/dev/kbd", O_RDONLY)) == -1)
    {
        fprintf(stderr, "Error with open() on /dev/kbd. Make sure exists.\n");
        perror (NULL);
        exit(EXIT_FAILURE);
    }

    while(1)
    {
        switch(toggle)
        {
            case 1:
            {
                /*
                 * Turn on Num Lock and make sure that
                 * Caps and Scroll lock are turned off.
                 */
                data = (_CONCTL_NUM_CHG | _CONCTL_NUM) | _CONCTL_CAPS_CHG | _CONCTL_SCROLL_CHG;
                break;
            }
            case 2:
            {
                /*
                 * Turn off Num Lock and now turn on Caps Lock
                 * (Scroll lock is already off).
                 */
                data = _CONCTL_NUM_CHG | (_CONCTL_CAPS_CHG | _CONCTL_CAPS);
                break;
            }
            case 3:
            {
                /*
                 * Turn off Caps lock and turn on Scroll lock
                 * (Num lock is already off).
                 */
                data = _CONCTL_CAPS_CHG | (_CONCTL_SCROLL_CHG | _CONCTL_SCROLL);
                toggle = 0;
                break;
            }
        }

        /* Explanation below. */
        if (error = devctl (fd, DCMD_CHR_SERCTL, &data,
                          sizeof(data), NULL))
        {
            fprintf(stderr, "Error setting KBD: %s\n",
                    strerror (error ));
            exit(EXIT_FAILURE);
        }
    }
}

```

```
    sleep(1);
    toggle++;
}

return (1);
}
```

Here's a quick explanation of the above *devctl()* call:

```
devctl (fd, DCMD_CHR_SERCTL, &data, sizeof(data), NULL)
```

The first parameter, *fd*, is the file descriptor of the device that's being changed. The second parameter is the device class that's being changed. In this case, it's a character device DCMD_CHR, with a "subclass" of _SERCTL. The third parameter is the data variable; this is the ORed value.

Example 3: Duration example

In this code, *tcdropline()*, which is used to disconnect a communications line, uses *devctl()* (this is the actual source code, *tcdropline()* is a standard library function):

```
#include <termios.h>
#include <devctl.h>
#include <errno.h>
#include <sys/dcmd_chr.h>

int tcdropline(int fd, int duration) {
    int error;

    duration = ((duration ? duration : 300) << 16) |
        _SERCTL_DTR_CHG | 0;

    if(error = devctl(fd, DCMD_CHR_SERCTL, &duration, sizeof duration, 0) == -1) {
        if(error == ENOSYS) {
            errno = ENOTTY;
        }
        return -1;
    }
    return 0;
}
```


Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

When *devctl()* fails, the effect of the failed command depends on the device driver. The corresponding data might be transferred, partially transferred, or not transferred at all.

The *devctl()* function was originally part of the POSIX 1003.1d draft standard; but it was deprecated in the IEEE *Approved Draft 10* standard.

See also:

close(), *open()*, *read()*, *write()*

difftime()

© 2004, QNX Software Systems Ltd.

Calculate the difference between two times

Synopsis:

```
#include <time.h>

double difftime( time_t time1,
                 time_t time0 );
```

Arguments:

time1, time0 The times to compare, expressed as **time_t** objects.

Library:

libc

Description:

The *difftime()* function calculates the difference between the calendar times specified by *time1* and *time0*:

time1 - time0

Returns:

The difference between the two times (in seconds).

Examples:

```
#include <stdio.h>
#include <time.h>
#include <stdlib.h>

void compute( void )
{
    int i, j;

    for( i = 1; i <= 20; i++ ) {
        for( j = 1; j <= 20; j++ ) {
            printf( "%3d ", i * j );
        }
        printf( "\n" );
    }
}
```

```
int main( void )
{
    time_t start_time, end_time;

    start_time = time( NULL );
    compute();
    end_time = time( NULL );
    printf( "Elapsed time: %f seconds\n",
           difftime( end_time, start_time ) );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

asctime(), *clock()*, *ctime()*, *gmtime()*, *localtime()*, *mktime()*, *strftime()*, *time()*, *tzset()*

dirctl()

© 2004, QNX Software Systems Ltd.

Control an open directory

Synopsis:

```
#include <dirent.h>

int dirctl( DIR * dir,
           int cmd,
           ... );
```

Arguments:

- dir* Provide control for this directory.
- cmd* At least the following values are defined in `<dirent.h>`:
- D_GETFLAG — retrieve the flags associated with the directory referenced by *dir*. For more information, see “Flag values,” below.
 - D_SETFLAG — set the flags associated with the directory referenced by *dir* to the value given as an additional argument. The new value can be any combination of the flags described in “Flag values,” below.

Library:

libc

Description:

The *dirctl()* function provides control over the open directory referenced by the *dir* argument. This function behaves in a manner similar to the file control function, *fcntl()*.

Flag values

D_FLAG_FILTER

Filter out duplicate name entries that may occur due to the union filesystem during a *readdir()* operation.

D_FLAG_STAT

Indicate to servers that they should attempt to return extra *stat()* information as part of the *readdir()* operation.

Returns:

The return value depends on the value of *cmd*:

- D_GETFLAG** The flags associated with the directory, or -1 if an error occurs (*errno* is set).
- D_SETFLAG** 0 for success, or -1 if an error occurs (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <dirent.h>

int main(int argc, char **argv) {
    DIR *dp;
    int ret;

    if(!(dp = opendir("/"))) {
        exit(EXIT_FAILURE);
    }

    /* Display the flags that are set on the
       directory by default*/
    if((ret = dircntl(dp, D_GETFLAG)) == -1) {
        exit(EXIT_FAILURE);
    }

    if(ret & D_FLAG_FILTER) {
        printf("Directory names are filtered\n");
    } else {
        printf("Directory names are not filtered\n");
    }

    if(ret & D_FLAG_STAT) {
        printf("Servers asked for extra stat information\n");
    } else {
        printf("Servers not asked for extra stat information\n");
    }

    closedir(dp);

    return 0;
}
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fcntl(), opendir()

Synopsis:

```
#include <libgen.h>

char *dirname( char *path );
```

Arguments:

path The string to parse.

Library:

`libc`

Description:

The *dirname()* function takes a pointer to a character string that contains a pathname, and returns a pointer to a string that's a pathname of the parent directory of that file. Trailing "/" characters in the path aren't counted as part of the path.

If the path *doesn't* contain a "/" character, or *path* is a NULL pointer or points to an empty string, then *dirname()* function returns a pointer to the string "." (dot).

Together the *dirname()* and *basename()* functions yield a complete pathname. The expression *dirname(path)* obtains the pathname of the directory where *basename(path)* is found.

Returns:

A pointer to a string that's the parent directory of *path*. If *path* is a NULL pointer or points to an empty string, a pointer to a string "." is returned.

Examples:

String input	String output
---------------------	----------------------

“/usr/lib”	“/usr”
------------	--------

“/usr/”	“usr”
---------	-------

“/”	“/”
-----	-----

“.”	“.”
-----	-----

“..”	“.”
------	-----

The following code fragment reads a pathname, changes the current working directory to the parent directory, and opens the file:

```
char path[BUFF_SIZE], *pathcopy;
int fd;

fgets(path, BUFF_SIZE, stdin);
pathcopy = strdup(path);
chdir(dirname(pathcopy));
fd = open(basename(path), O_RDONLY);
```

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

basename()

dispatch_block()

© 2004, QNX Software Systems Ltd.

Block while waiting for an event

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

dispatch_context_t * dispatch_block
    ( dispatch_context_t * ctp );
```

Arguments:

ctp A pointer to a `dispatch_context_t` structure that defines the dispatch context.

Library:

`libc`

Description:

The `dispatch_block()` function blocks while waiting for an event (e.g. message or signal) that's registered using one of the attach functions, `message_attach()`, `pulse_attach()`, `resmgr_attach()`, or `select_attach()`. (The `sigwait_attach()` function isn't currently implemented.)

If the type of blocking is:	<i>dispatch_block()</i> does a:
message (resmgr, message, select)	<i>MsgReceive()</i>
signal	<i>SignalWaitinfo()</i>

Returns:

A dispatch context that's passed in by `dispatch_context_alloc()`. or NULL if an error occurs (`errno` is set).

Errors can occur when the blocking kernel call returns with an error, for example, due to the delivery of a signal.



In the case of a timeout, a valid *ctp* is returned, but either the *ctp->message_context.rcvid* or *ctp->sigwait_context.signo* is set to -1.

If a non-NULL context pointer is returned, it could be different from the one passed in, as it's possible for the *ctp* to be reallocated to a larger size. In this case, the old *ctp* is no longer valid. However, if NULL is returned (for example, because a signal interrupted the *MsgReceive()*), the old context pointer is still valid. Typically, a resource manager would target signals to a thread dedicated to handling signals. However, if a signal can be targeted to the thread doing *dispatch_block()*, you could use the following code in this situation:

```
dispatch_context_t  *ctp, *new_ctp;

ctp = dispatch_context_alloc( ... );
while (1) {
    new_ctp = dispatch_block( ctp );
    if ( new_ctp ) {
        ctp = new_ctp
    }
    else {
        /* handle the error condition */
        :
    }
}
```

Errors:

EFAULT	A fault occurred when the kernel tried to access the buffers.
EINTR	The call was interrupted by a signal.
EINVAL	Invalid arguments passed to <i>dispatch_block()</i> .
ENOMEM	Insufficient memory to allocate internal data structures.

See also the error constants returned in *MsgReceive()* and *SignalWaitinfo()*.

Examples:

```
#include <sys/dispatch.h>

int main( int argc, char **argv ) {
    dispatch_context_t *ctp;

    :

    for(;;) {
        if( ctp = dispatch_block( ctp ) ) {
            dispatch_handler( ctp );
        }
    }
}
```

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

dispatch_context_alloc(), *dispatch_handler()*, *dispatch_timeout()*, *dispatch_unblock()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

dispatch_context_t * dispatch_context_alloc
( dispatch_t * dpp );
```

Arguments:

dpp A dispatch handle created by *dispatch_create()*.

Library:

libc

Description:

The *dispatch_context_alloc()* function returns a dispatch context pointer. The function is passed in the handle *dpp* from *dispatch_create()*. The dispatch context is used by dispatch to do its work. It's passed as an argument to *dispatch_block()* and *dispatch_handler()*.



The *dispatch_context_alloc()* function fails if you haven't attached any events to dispatch yet (e.g. you didn't call *message_attach()*, *resmgr_attach()*, or *select_attach()*). The dispatch library can't allocate a proper context until it knows what kind of events you want to block.

Returns:

A pointer to a dispatch context, or NULL if an error occurs (*errno* is set).

Errors:

- EINVAL No events were attached.
- ENOMEM Insufficient memory to allocate context.

Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char **argv ) {
    dispatch_t      *dpp;
    dispatch_context_t *ctp;

    if( ( dpp = dispatch_create() ) == NULL ) {
        fprintf( stderr, "%s: Unable to allocate \
            dispatch handle.\n", argv[0] );
        return EXIT_FAILURE;
    }

    &vellip;

    ctp = dispatch_context_alloc( dpp );

    &vellip;

    return EXIT_SUCCESS;
}
```

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

- Cancellation point No
- Interrupt handler No
- Signal handler No

continued...

Safety

Thread	Yes
--------	-----

See also:

dispatch_block(), *dispatch_context_free()*, *dispatch_create()*,
dispatch_handler(), *dispatch_unblock()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

dispatch_context_free()

© 2004, QNX Software Systems Ltd.

Free a dispatch context

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

void dispatch_context_free(
    dispatch_context_t * ctp );
```

Arguments:

ctp A pointer to a `dispatch_context_t` structure that was allocated by `dispatch_context_alloc()`.

Library:

libc

Description:

The `dispatch_context_free()` function frees the given dispatch context.

Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char **argv ) {
    dispatch_t      *dpp;
    dispatch_context_t *ctp;

    if( ( dpp = dispatch_create() ) == NULL ) {
        fprintf( stderr, "%s: Unable to allocate
            dispatch handle.\n", argv[0] );
        return EXIT_FAILURE;
    }

    :

    ctp = dispatch_context_alloc( dpp );

    :

    dispatch_context_free ( ctp );
```



```
    return EXIT_SUCCESS;  
}
```

See *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()* for examples using the dispatch interface.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

dispatch_context_alloc()

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer's Guide*.

dispatch_create()

© 2004, QNX Software Systems Ltd.

Allocate a dispatch handle

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

dispatch_t *dispatch_create( void );
```

Library:

libc

Description:

The *dispatch_create()* function allocates and initializes a dispatch handle. The attach functions are:

- *message_attach()*
- *pulse_attach()*
- *resmgr_attach()*
- *select_attach()*

If you wish, you can do a *resmgr_attach()* with a NULL path. This has the effect of initializing dispatch to receive messages and creates the channel among other things.



A channel is created only when you first attach something that requires a channel (indicating you will block receiving messages).

Returns:

A handle to a dispatch structure, or NULL if an error occurs.



The dispatch structure, **dispatch_t**, is an opaque data type; you can't access its contents directly.

Errors:

ENOMEM Insufficient memory to allocate context.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <stddef.h>
#include <fcntl.h>
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int my_func( select_context_t *ctp, int fd,
            unsigned flags, void *handle ) {
    int  i, c;

    /* Do some useful stuff with data */
    i = read( fd, &c, 1 );
    fprintf( stderr, "activity on fd %d: read char %c,
                  return code %d\n", fd, c, i );
}

int main( int argc, char **argv ) {
    dispatch_t          *dpp;
    dispatch_context_t  *ctp;
    select_attr_t        attr;
    int                  fd, fd2;

    if( ( dpp = dispatch_create() ) == NULL ) {
        fprintf( stderr, "%s: Unable to allocate \
                  dispatch handle.\n", argv[0] );
        return EXIT_FAILURE;
    }

    if( argc ≤ 2 || (fd = open( argv[1],
                               O_RDWR | O_NONBLOCK )) == -1 ) {
        return EXIT_FAILURE;
    }

    if( argc ≤ 2 || (fd2 = open( argv[2],
                                 O_RDWR | O_NONBLOCK )) == -1 ) {
        return EXIT_FAILURE;
    }

    select_attach( dpp, &attr, fd,
                  SELECT_FLAG_READ | SELECT_FLAG_REARM, my_func, NULL );
    select_attach( dpp, &attr, fd2,
                  SELECT_FLAG_READ | SELECT_FLAG_REARM, my_func, NULL );
    ctp = dispatch_context_alloc( dpp );
}
```

```
for(;;) {
    if( ctp = dispatch_block( ctp ) ) {
        dispatch_handler( ctp );
    }
}
return EXIT_SUCCESS;
}
```

For more examples using the dispatch interface, see *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

dispatch_block(), *dispatch_context_alloc()*, *dispatch_destroy()*,
dispatch_handler(), *dispatch_timeout()*, *dispatch_unblock()*
message_attach(), *pulse_attach()*, *resmgr_attach()*, *select_attach()*

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int dispatch_destroy( dispatch_t *dpp );
```

Arguments:

dpp A dispatch handle created by *dispatch_create()*.

Library:

libc

Description:

The function *dispatch_destroy()* destroys the given dispatch handle.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EINVAL The dispatch handle, *dpp*, is invalid.

Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char **argv ) {
    dispatch_t    *dpp;
    int           destroyed;

    if( ( dpp = dispatch_create() ) == NULL ) {
        fprintf( stderr, "%s: Unable to allocate \
                dispatch handle.\n", argv[0] );
        return EXIT_FAILURE;
    }
}
```

```
    }  
    :  
    if ( (destroyed = dispatch_destroy ( dpp )) == -1 ) {  
        fprintf ( stderr, "Dispatch wasn't destroyed, \  
            bad dispatch handle %d.\n", dpp);  
        return EXIT_FAILURE;  
    }  
    /* else dispatch was destroyed */  
    :  
    return EXIT_SUCCESS;  
}
```

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

dispatch_create()

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int dispatch_handler( dispatch_context_t * ctp );
```

Arguments:

ctp A pointer to a `dispatch_context_t` structure that was allocated by `dispatch_context_alloc()`.

Library:

`libc`

Description:

The `dispatch_handler()` function handles events received by `dispatch_block()`. Depending on the blocking type, `dispatch_handler()` does one of the following:

- Calls the `message_*` subsystem. A search is made (based upon the message type or pulse code) for a matching function (that was attached with `message_attach()` or `pulse_attach()`). If a match is found, the attached function is called.
- If the message type is in the range handled by the resource manager (e.g. I/O messages) and pathnames were attached using `resmgr_attach()`, then the `resmgr_*` subsystem is called and handles the resource manager messages.
- If a pulse is received, it may be dispatched to the `resmgr_*` subsystem if it's one of the codes (unblock and disconnect pulses) handled by the resource manager. If a `select_attach()` was done and the pulse matches the one used by `select_attach()`, then the `select_*` subsystem is called and dispatches that event.
- If a message is received, and no matching handler is found for that message type, `MsgError()` returns ENOSYS to the sender.

- If a *SignalWaitinfo()* blocking type is used, then a search is made based upon the signal number for a matching function attached by the program (using the *sigwait_attach()* function, not currently implemented). If a match is found, that function is called.

Returns:

- 0 Success.
- 1 One of the following occurred:
 - The message was a `_PULSE_CODE_THREADDEATH` pulse message (see *ChannelCreate()*) for which there's no default handler function.
 - The message length was less than 2 bytes. A 2-byte message type is required at the beginning of the message so that a handler function can be found or identified.
 - The message wasn't in native endian format and there were no handler functions that specified `MSG_FLAG_CROSS_ENDIAN` on this range, even though a handler was registered for it using *message_attach()*. The `MSG_FLAG_CROSS_ENDIAN` flag wasn't given to *message_attach()*.
 - A handler was found for the message, but the handler determined that there was a problem.

In any case, if the message wasn't a pulse, then the client will be replied to with an appropriate *errno*.

Examples:

```
#include <stdlib.h>
#include <sys/dispatch.h>

int main( int argc, char **argv ) {
    dispatch_context_t *ctp;

    :

    for(;;) {
        if( ctp = dispatch_block( ctp ) ) {
```



```
        dispatch_handler( ctp );
    }
}
return EXIT_SUCCESS;
}
```

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point	Read the <i>Caveats</i>
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

This function might or might not be a cancellation point, depending on the implementation of the handler.

See also:

dispatch_block(), *dispatch_create()*, *dispatch_timeout()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

dispatch_timeout()

© 2004, QNX Software Systems Ltd.

Set a timeout

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int dispatch_timeout( dispatch_t *dpp,
                    struct timespec *reltime );
```

Arguments:

dpp A dispatch handle created by *dispatch_create()*.

reltime A pointer to a **timespec** structure that specifies the relative time of the timeout.

Library:

libc

Description:

The function *dispatch_timeout()* sets a timeout that's used when blocking with *dispatch_block()*.

Returns:

0 Success.

-1 An error occurred.

Examples:

```
#include <sys/dispatch.h>
#include <time.h>
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char **argv ) {
    dispatch_t        *dpp;
    struct timespec    time_out;
    int                timedout;
    time_out.tv_sec = 1;
    time_out.tv_nsec = 2;
```

```

if( ( dpp = dispatch_create() ) == NULL ) {
    fprintf( stderr, "%s: Unable to allocate \
            dispatch handle.\n", argv[0] );
    return EXIT_FAILURE;
}

:

if ( (timeout = dispatch_timeout ( dpp, &time_out )
    == -1 ) {
    fprintf ( stderr, "Couldn't set timeout );
    return EXIT_FAILURE;
}
/* else successful timeout set */

:
return EXIT_SUCCESS;
}

```

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

dispatch_block(), *dispatch_create()*, *dispatch_handler()*,
dispatch_unblock() `timespec`

dispatch_unblock()

© 2004, QNX Software Systems Ltd.

Unblock all of the threads that are blocked on a dispatch handle

Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

void dispatch_unblock( dispatch_context_t * ctp );
```

Arguments:

ctp A pointer to a `dispatch_context_t` structure that defines the dispatch context.

Library:

`libc`

Description:

This routine tries to unblock all of the threads that are blocked on the given dispatch handle. You should use this function in the thread pool structure as the unblock function pointer so that *thread_pool_control()* will behave properly.

Currently, this function unblocks only channel resources.

Examples:

For examples using the dispatch interface, see *dispatch_create()*, *message_attach()*, *resmgr_attach()*, and *thread_pool_create()*.

Classification:

QNX Neutrino

Safety

Cancellation point Yes

Interrupt handler No

continued...

Safety

Signal handler	Yes
Thread	Yes

See also:

dispatch_block(), *dispatch_context_alloc()*, *dispatch_handler()*,
dispatch_timeout()

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

div()

© 2004, QNX Software Systems Ltd.

Calculate a quotient and remainder

Synopsis:

```
#include <stdlib.h>

div_t div( int numer,
           int denom );
```

Arguments:

numer The numerator in the division.
denom The denominator.

Library:

libc

Description:

The *div()* function calculates the quotient and remainder of the division of *numer* by *denom*.

Returns:

A `div_t` structure containing the quotient and remainder:

```
typedef struct {
    int quot;      /* quotient */
    int rem;       /* remainder */
} div_t;
```

Examples:

```
#include <stdio.h>
#include <stdlib.h>

void print_time( int seconds )
{
    div_t min_sec;

    min_sec = div( seconds, 60 );
    printf( "It took %d minutes and %d seconds\n",
```

```
        min_sec.quot, min_sec.rem );
    }

int main( void )
{
    print_time( 130 );

    return EXIT_SUCCESS;
}
```

produces the output:

```
It took 2 minutes and 10 seconds
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

ldiv()

dladdr()

© 2004, QNX Software Systems Ltd.

Translate an address to symbolic information

Synopsis:

```
#include <dlfcn.h>

int dladdr( void *address,
            Dl_info *dli );
```

Arguments:

address The address for which you want symbolic information.

dli A pointer to a **Dl_info** structure where the function can store the symbolic information. Your application must allocate the space for this structure; *dladdr()* fills in the members, based on the specified address.

The **Dl_info** structure includes the following members:

- **const char *dli_fname** — a pointer to the filename of the object containing *address*.
- **void *dli_fbase** — the base address of the object containing *address*.
- **const char *dli_sname** — a pointer to the symbol name nearest the specified *address*. This symbol is either at *address*, or is the nearest symbol with a lower address.
- **void *dli_saddr** — the actual address of the *dli_sname* symbol.

If *dladdr()* can't find a symbol that describes the specified *address*, the function sets *dli_sname* and *dli_saddr* to NULL.

Library:

libc

Description:

The *dladdr()* function determines whether the specified *address* is located within one of the objects that make up the calling application's address space.



The *dladdr()* function is available only to dynamically linked processes.

Returns:

0 if the specified *address* can't be matched, or nonzero if it could be matched.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

The `DL_info` pointers may become invalid if objects are removed via *dlclose()*.

There's no way to determine which symbol you'll get if multiple symbols are mapped to the same address.

See also:

dlclose(), dlerror(), dlopen(), dlsym()

Synopsis:

```
#include <dlfcn.h>

int dlclose( void *handle );
```

Arguments:

handle A handle for a shared object, returned by *dlopen()*.

Library:

`libc`

Description:

The *dlclose()* function disassociates a shared object opened by *dlopen()* from the calling process. An object's symbols are no longer available after it's been closed with *dlclose()*. All objects loaded as a result of the closed objects dependencies are also closed.

The *handle* argument is the value returned by a previous call to *dlopen()*.



The *dlclose()* function is available only to dynamically linked processes.

Returns:

0 for success, or a nonzero value if an error occurs.

Errors:

If an error occurs, more detailed diagnostic information is available from *dlerror()*.

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

An object won't be removed from the address space until all references to that object (via *dlopen()* or dependencies from other objects) have been closed.

Referencing a symbol in a closed object can cause undefined behavior.

See also:

dladdr(), *dlerror()*, *dlopen()*, *dlsym()*

Synopsis:

```
#include <dldfcn.h>

char *dlderror( void );
```

Library:

libc

Description:

The *dlderror()* function returns a NULL-terminated string (with no trailing newline) describing the last error that occurred during a call to one of the *dld*()* functions. If no errors have occurred, *dlderror()* returns NULL.



The *dldopen()* function is available only to dynamically linked processes.

Returns:

A pointer to an error description, or NULL.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

dladdr(), dlclose(), dlopen(), dlsym()

Synopsis:

```
#include <dlfcn.h>

void * dlopen( const char * pathname,
               int mode );
```

Arguments:

pathname NULL, or the path to the executable object file that you want to access.

mode Flags that control how *dlopen()* operates; see “The *mode*,” below.

Library:

`libc`

Description:

The *dlopen()* function gives you direct access to the dynamic linking facilities by making the executable object file specified in *pathname* available to the calling process. It returns a handle that you can use in subsequent calls to *dlsym()* and *dlclose()*.



The *dlopen()* function is available only to a dynamically-linked process. A statically-linked process (one where `libc` is linked statically) can't call *dlopen()* because a statically-linked executable:

- doesn't export any of its symbols
- can't export the required structure for libraries to link against
- can't fill structures at startup needed to load subsequent shared objects.

Any dependencies recorded within *pathname* are loaded as part of the *dlopen()* call. These dependencies are searched in load-order to locate

any additional dependencies. This process continues until all of the dependencies for *pathname* have been satisfied. This dependency tree is called a *group*.

If *pathname* is NULL, *dlopen()* provides a handle to the running process's global symbol object. This provides access to the symbols from the original program image file, the dependencies it loaded at startup, plus any objects opened with *dlopen()* calls using the RTLD_GLOBAL flag. This set of symbols can change dynamically if the application subsequently calls *dlopen()* using RTLD_GLOBAL.

You can use *dlopen()* any number of times to open objects whose names resolve to the same absolute or relative path name; the object is loaded into the process's address space only once.

In order to find the shared objects, the following directories or paths are searched in order:

- **RPATH**
- **LD_LIBRARY_PATH**
- **CS_LIBPATH.**

Note that **LD_LIBRARY_PATH** is ignored if the binary is **setuid**, and where the **euid** is not the same as the **uid** of the user running the binary. This is done for security purposes.



The above directories are set as follows:

- The **RPATH** value is set up when binary is linked, using the **ld** command line option **-rpath**. See **ld** for details.
- The **LD_LIBRARY_PATH** is generally set up by other startup script, either in the boot image or by a secondary script. For example, on self hosted QNX system, it is setup by **ph** script. It is not part of any default environment.
- **CS_LIBPATH** is populated by the kernel, and the default value is based on the **LD_LIBRARY_PATH** value of the **procnto** command line in the boot image. Note that, you may use **getconf** utility to inspect this value and **setconf** to set this value. For example:

```
setconf CS_LIBPATH 'getconf CS_LIBPATH' :/new/path
```

When loading shared objects, the application should open a specific version instead of relying on the version pointed to by a symbolic link.

The *mode*

The *mode* argument indicates how *dlopen()* operates on *pathname* when handling relocations, and controls the visibility of symbols found in *pathname* and its dependencies.

The *mode* argument is a bitwise-OR of the constants described below. Note that the relocation and visibility constants are mutually exclusive.

Relocation

When you load an object by calling *dlopen()*, the object may contain references to symbols whose addresses aren't known until the object has been loaded; these references must be relocated before accessing the symbols. The *mode* controls when relocations take place, and can be one of:

- RTLD_LAZY References to data symbols are relocated when the object is loaded. References to functions aren't relocated until that function is invoked. This improves performance by preventing unnecessary relocations.

- RTLD_NOW All references are relocated when the object is loaded. This may waste cycles if relocations are performed for functions that never get called, but this behavior could be useful for applications that need to know that all symbols referenced during execution are available as soon as the object is loaded.



RTLD_LAZY isn't currently supported.

Visibility

The following *mode* bits determine the scope of visibility for symbols loaded with *dlopen()*:

RTLD_GLOBAL

Make the object's global symbols available to any other object. Symbol lookup using `dlopen(0, mode)` and an associated *dlsym()* are also able to find the object's symbols.

RTLD_LOCAL

Make the object's global symbols available only to objects in the same group.

The program's image and any objects loaded at program startup have a *mode* of RTLD_GLOBAL; the default *mode* for objects acquired with *dlopen()* is RTLD_LOCAL. A local object may be part of the dependencies for more than one group; any object with a RTLD_LOCAL *mode* referenced as a dependency of an object with a RTLD_GLOBAL *mode* is promoted to RTLD_GLOBAL.

Objects loaded with *dlopen()* that require relocations against global symbols can reference the symbols in any RTLD_GLOBAL object.

You can OR the *mode* with the following values to affect symbol scope:

RTLD_GROUP

Only symbols from the associated group are available. All dependencies between group members must be satisfied by the objects in the group.

RTLD_WORLD

Only symbols from RTLD_GLOBAL objects are available.

The default *mode* is RTLD_WORLD | RTLD_GROUP.



If you specify RTLD_WORLD without RTLD_GROUP, *dlopen()* doesn't load any of the DLL's dependencies.

Symbol Resolution

When resolving the symbols in the shared object, the runtime linker searches for them in the dynamic symbol table using the following order:

- By default:
- 1 main executable
 - 2 the shared object being loaded
 - 3 all other loaded shared objects that were loaded with the RTLD_GLOBAL flag.

When **-Bsymbolic** is specified:

- 1 the shared object being loaded
- 2 main executable
- 3 all other loaded shared objects that were loaded with the RTLD_GLOBAL flag.

For executables, the dynamic symbol table typically contains only those symbols that are known to be needed by any shared libraries.

This is determined by the linker when the executable is linked against a shared library.

Since you don't link your executable against a shared object that you load with *dlopen()*, the linker can't determine which executable symbols need to be made available to the shared object.

If your shared object needs to resolve symbols in the executable, then you may force the linker to make *all* of the symbols in the executable available for dynamic linking by specifying the **-E** linker option. For example:

```
gcc -Vgcc_ntox86 -Wl,-E -o main main.o
```

Shared objects always place all their symbols in dynamic symbol tables, so this option isn't needed when linking a shared object.

Returns:

A handle to the object, or NULL if an error occurs.



Don't interpret the value of this handle in any way. For example, if you open the same object repeatedly, don't assume that *dlopen()* returns the same handle.

Errors:

If an error occurs, more detailed diagnostic information is available from *dLError()*.

Environment variables:

DL_DEBUG Display debugging information about the libraries as they're opened.

LD_LIBRARY_PATH

The **LD_LIBRARY_PATH** environment variable is searched for any dependencies required by *pathname*.

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

Some symbols defined in executables or shared objects might not be available to the runtime linker. The symbol table created by `ld` for use by the runtime linker might contain a subset of the symbols defined in the object.

See also:

dladdr(), *dlclose()*, *dlderror()*, *dlsym()*

`ld`, `gcc` in the *Utilities Reference*

dlsym()

© 2004, QNX Software Systems Ltd.

Get the address of a symbol in a shared object

Synopsis:

```
#include <dlfcn.h>

void* dlsym( void* handle,
             const char* name );
```

Arguments:

handle Either a handle for a shared object, returned by *dlopen()*, or the special flag, *RTLD_DEFAULT*.

name The name of the symbol that you want to find in the shared object.

Library:

`libc`

Description:

The *dlsym()* function lets a process obtain the address of the symbol specified by *name* defined in a shared object.



The *dlsym()* function is available only to dynamically linked processes.

If *handle* is a handle returned by *dlopen()*, you must not have closed that shared object by calling *dlclose()*. The *dlsym()* functions also searches for the named symbol in the objects loaded as part of the dependencies for that object.

If *handle* is *RTLD_DEFAULT*, *dlsym()* searches all objects in the current process, in load-order.

In the case of *RTLD_DEFAULT*, if the objects being searched were loaded with *dlopen()*, *dlsym()* searches the object only if the caller is part of the same dependency hierarchy, or if the object was loaded with global search access (using the *RTLD_GLOBAL* mode).

Returns:

A pointer to the named symbol for success, or NULL if an error occurs.

Errors:

If an error occurs, more detailed diagnostic information is available from *dlerror()*.

Examples:

Use *dlsym()* to find a function pointer and a pointer to a global variable in a shared library:

```
typedef int (*foofunc)( int );

void* handle;
int* some_global_int;
foofunc brain;

/* Open a shared library. */
handle = dlopen( "/usr/nto/x86/lib/libfoo.so.1", RTLD_NOW );

/* Find the address of a function and a global integer. */
brain = (foofunc)dlsym( handle, "takeover_world" );
some_global_int = (int*)dlsym( handle, "my_global_int" );

/* Invoke the function and print the int. */
x = (*brain)( 5 );
printf( "that global is %d\n", *some_global_int );
```

Check to see if a function is defined, and call it if it is:

```
typedef int (*funcptr)( void );

funcptr funk = NULL;

funk = (funcptr)dlsym( RTLD_DEFAULT, "get_funky" );
if( funk != NULL ) {
    (*funk)();
}
```

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

Function pointers are a pain; use **typedefs** to help preserve your sanity.

See also:

dladdr(), dlclose(), dlerror(), dlopen()

Synopsis:

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int dn_comp( const char * exp_dn,
             u_char * comp_dn,
             int length,
             u_char ** dnptrs,
             u_char ** lastdnptr );
```

Arguments:

exp_dn The Internet domain name you want to compress.

comp_dn A buffer where the function can store the compressed name.

length The size of the array that *comp_dn* points to.

dnptrs NULL, or an array of pointers to previously compressed names in the current message; see below.

lastdnptr NULL, or the limit of the array specified by *dnptrs*.

Library:

`libsocket`

Description:

The *dn_comp()* routine is a low-level routine used by *res_query()* to compress an Internet domain name. This routine compresses the domain name *exp_dn* and stores it in *comp_dn*.

The compression uses an array of pointers, *dnptrs*, to previously compressed names in the current message. The first pointer points to the beginning of the message and the list ends with NULL. The limit to the array is specified by *lastdnptr*. As a side effect, *dn_comp()*

updates the list of pointers for labels inserted into the message as the name is compressed. If *dnptrs* is NULL, names aren't compressed. If *lastdnptr* is NULL, the list of labels isn't updated.

Returns:

The size of the compressed domain name, in bytes, or -1 if an error occurs.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

dn_expand(), *gethostbyname()*, *res_init()*, *res_mkquery()*, *res_query()*, *res_search()*, *res_send()*

/etc/resolv.conf, *hostname* in the *Utilities Reference*

RFC 974, *RFC 1032*, *RFC 1033*, *RFC 1034*, *RFC 1035*

Synopsis:

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int dn_expand( const u_char * msg,
              const u_char * eomorig,
              const u_char * comp_dn,
              char * exp_dn,
              int length );
```

Arguments:

msg A pointer to the beginning of the message that contains the compressed name.

eomorig A pointer to the first location after the message.

comp_dn The compressed name that you want to expand.

exp_dn A buffer where the function can store the expanded name.

length The length of the array specified by *exp_dn*.

Library:

`libsocket`

Description:

The *dn_expand()* function is a low-level routine used by *res_query()* to expand the compressed domain name, *comp_dn*, to a full domain name.

The compressed name is contained in a query or reply message.

Returns:

The size of the *compressed* domain name (not the expanded name), in bytes, or -1 if an error occurs.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

dn_comp(), *gethostbyname()*, *res_init()*, *res_mkquery()*, *res_query()*,
res_search(), *res_send()*

/etc/resolv.conf, *hostname* in the *Utilities Reference*

RFC 974, *RFC 1032*, *RFC 1033*, *RFC 1034*, *RFC 1035*

Synopsis:

```
#include <stdlib.h>

double drand48( void );
```

Library:

libc

Description:

The *drand48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a nonnegative **double** uniformly distributed over the interval [0.0, 1.0].

Call one of *lcg48()*, *seed48()*, or *srand48()* to initialize the random-number generator before calling *drand48()*, *lrand48()*, or *mrnd48()*.

The *erand48()* function is a thread-safe version of *drand48()*.

Returns:

A pseudo-random **double**.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	No

See also:

*erand48(), jrand48(), lcong48(), lrand48(), mrand48(), nrand48(),
seed48(), srand48()*

Synopsis:

```
#include <math.h>

double drem ( double x,
              double y );

float dremf ( float x,
             float y );
```

Arguments:

- x* The numerator of the division.
- y* The denominator.

Library:

`libm`

Description:

The *drem()* and *dremf()* functions compute the remainder $r = x - n * y$, when *y* is nonzero. The value *n* is the integral value nearest the exact value x/y .

When $|n - x/y| = \frac{1}{2}$, the value *n* is chosen to be even. But *remainder(x, 0)* and *remainder(infinity,0)* are invalid operations that produce a NAN.

The behavior of *drem()* is independent of the rounding mode.

Returns:

The remainder, $r = x - n * y$, when *y* is nonzero.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

remainder()

Synopsis:

```
#include <ds.h>

int ds_clear( ds_t dsdes,
              const char* variable_name );
```

Arguments:

dsdes A data server descriptor returned by *ds_register()*.
variable_name The name of the variable that you want to delete.

Library:

`libds`

Description:

The *ds_clear()* function deletes *variable_name* from the data server identified by *dsdes*.

Returns:

0 Success.
-1 An error occurred (*errno* is set).

Errors:

EBADF Invalid file descriptor *dsdes*.
ESRCH The variable doesn't exist in the data server.

Examples:

See `slinger` in the *Utilities Reference*.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_deregister(), *ds_flags()*

Synopsis:

```
#include <ds.h>

int ds_create( ds_t dsdes,
              const char * variable_name,
              char flags,
              struct sigevent * sigevent );
```

Arguments:

<i>dsdes</i>	A data server descriptor returned by <i>ds_register()</i> .
<i>variable_name</i>	The name of the variable that you want to create. All variables are global, so only one instance of the variable can exist in the data server process. The maximum length of a variable name is 60 characters.
<i>flags</i>	Flags that specify the variable's behavior: <ul style="list-style-type: none">• DS_PERM — don't delete the variable when the application that created it terminates. The variable is removed when the data server process terminates, or if the flag is turned off after the application that created the variable terminates. If <i>flags</i> is 0, the variable is removed if you call <i>ds_deregister()</i> , or the process terminates.
<i>sigevent</i>	A pointer to a sigevent structure that describes a proxy or signal to be sent to the external application that created the variable if the data referenced by the variable changes; see below.

Library:

libds

Description:

The *ds_create()* function creates a variable, whose name is given by *variable_name*, on the data server identified by *dsdes*.

If the data referenced by *variable_name* changes, a proxy or signal, described in the **sigevent** structure, can be sent to the external application that created *variable_name* (see *ds_set()*).

We recommend the following event types for use with this function:

- SIGEV_SIGNAL
- SIGEV_SIGNAL_CODE
- SIGEV_SIGNAL_THREAD
- SIGEV_PULSE
- SIGEV_INTR

To display the current value of a variable on an HTML page, use the **qnxvar** token with the **read** tag. See the description of **slinger** in the *Utilities Reference*.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- | | |
|--------|--|
| EBADF | Invalid file descriptor <i>dsdes</i> . |
| EEXIST | The variable name already exists in the data server. |
| ENOMEM | Not enough memory to create the variable or initialize the data. |

Examples:

See `slinger` in the *Utilities Reference*.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_flags(), *ds_get()*, *ds_register()*, *ds_set()*, `sigevent`

ds_deregister()

© 2004, QNX Software Systems Ltd.

Deregister an application with the data server

Synopsis:

```
#include <ds.h>

int ds_deregister( ds_t dsdes );
```

Arguments:

dsdes A data server descriptor returned by *ds_register()*.

Library:

`libds`

Description:

The *ds_deregister()* function deregisters your application with the data server, *dsdes*, and deletes any variables that the data server process created, except those with the DS_PERM flag set (see *ds_flags()*).

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EBADF Invalid file descriptor, *dsdes*.

Examples:

See `slinger` in the *Utilities Reference*.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_flags(), *ds_register()*

ds_flags()

© 2004, QNX Software Systems Ltd.

Set the flags for a data server variable

Synopsis:

```
#include <ds.h>

int ds_flags( ds_t dsdes,
              const char* variable_name,
              char flags );
```

Arguments:

<i>dsdes</i>	A data server descriptor returned by <i>ds_register()</i> .
<i>variable_name</i>	The name of the data server variable.
<i>flags</i>	The new flags for the variable. The flags include: <ul style="list-style-type: none">• DS_PERM — don't delete the variable when the application that created it terminates. The variable is removed when the data server process terminates, or if the flag is turned off after the application that created the variable terminates.

Library:

`libds`

Description:

The *ds_flags()* function changes the state of the *flags* belonging to the variable called *variable_name* on the data server identified by *dsdes*.

Returns:

0 for success, or -1 if an error occurs (*errno* is set).

Errors:

- EBADF Invalid file descriptor *dsdes*.
- ESRCH The variable doesn't exist in the data server.

Classification:

QNX Neutrino

Safety

- Cancellation point Yes
- Interrupt handler No
- Signal handler Yes
- Thread Yes

See also:

ds_clear(), *ds_create()*, *ds_deregister()*, *ds_set()*

ds_get()

© 2004, QNX Software Systems Ltd.

Retrieve a data server variable

Synopsis:

```
#include <ds.h>

int ds_get( ds_t dsdes,
            const char* variable_name,
            const char* variable_data,
            size_t data_len );
```

Arguments:

<i>dsdes</i>	A data server descriptor returned by <i>ds_register()</i> .
<i>variable_name</i>	The name of the data server variable that you want to get.
<i>variable_data</i>	A buffer where the function can store the data associated with the variable.
<i>data_len</i>	The size of the buffer, in bytes.

Library:

`libds`

Description:

The *ds_get()* function retrieves the data corresponding to *variable_name* from the data server *dsdes*, and places it in the buffer pointed to by *variable_data*.

Returns:

The amount of data written to the buffer *variable_data*, or -1 if an error occurs (*errno* is set).

Errors:

EBADF	Invalid file descriptor <i>dsdes</i> .
EMSGSIZE	The buffer isn't big enough for the data.
ESRCH	The variable doesn't exist in the data server.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_create(), *ds_set()*

ds_register()

© 2004, QNX Software Systems Ltd.

Register an application with the data server

Synopsis:

```
#include <ds.h>

ds_t ds_register( void );
```

Library:

libds

Description:

The *ds_register()* function registers your application with the data server. The data server must reside on the same node as your application.

Returns:

A data server descriptor, or -1 if an error occurs (*errno* is set).

Errors:

ENOENT	No such file or directory; the data server isn't started.
ENOMEM	Insufficient memory is available to communicate with the data server.

Examples:

See **slinger** in the *Utilities Reference*.

Classification:

QNX Neutrino

Safety

Cancellation point Yes

continued...

Safety

Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_deregister()

ds_set()

© 2004, QNX Software Systems Ltd.

Set a data server variable

Synopsis:

```
#include <ds.h>

int ds_set( ds_t dsdes,
            const char* variable_name,
            const char* variable_data,
            size_t data_len );
```

Arguments:

<i>dsdes</i>	A data server descriptor returned by <i>ds_register()</i> .
<i>variable_name</i>	The name of the data server variable that you want to set.
<i>variable_data</i>	A pointer to the data you want to associate with the variable.
<i>data_len</i>	The size of the data, in bytes.

Library:

libds

Description:

The *ds_set()* function passes the data *variable_data* to the data server identified by *dsdes*. The data server stores the data in the variable whose name is given by *variable_name*, overwriting any existing value.

To display the modified data on an HTML page, use the **qnxvar** token with the **read** tag. See the description of **slinger** in the *Utilities Reference*.

Returns:

0 for success, or -1 if an error occurs (*errno* is set).

Errors:

EBADF Invalid file descriptor *dsdes*.
ENOMEM Not enough memory to store the data.
ESRCH The variable doesn't exist in the data server.

Examples:

See **slinger** in the *Utilities Reference*.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

ds_create(), *ds_flags()*, *ds_get()*

dup()

© 2004, QNX Software Systems Ltd.

Duplicate a file descriptor

Synopsis:

```
#include <unistd.h>

int dup( int filedes );
```

Arguments:

filedes The file descriptor that you want to duplicate.

Library:

`libc`

Description:

The *dup()* function duplicates the file descriptor specified by *filedes*. The new file descriptor refers to the same open file descriptor as the original, and shares any locks. The new file descriptor also:

- references the same file or device
- has the same open mode (read and/or write)
- has an identical file position to the original (changing the position with one descriptor results in a changed position in the other).

Changing the file position with one descriptor results in a changed position for the other.

Calling:

```
dup_filedes = dup( filedes );
```

is the same as:

```
dup_filedes = fcntl( filedes, F_DUPFD, 0 );
```


Returns:

The new file descriptor for success, or -1 if an error occurs (*errno* is set).

Errors:

- | | |
|--------|--|
| EBADF | The file descriptor, <i>filedes</i> , isn't a valid. |
| EMFILE | There are already OPEN_MAX file descriptors in use. |
| ENOSYS | The <i>dup()</i> function isn't implemented for the filesystem specified by <i>filedes</i> . |

Examples:

```
#include <fcntl.h>
#include <unistd.h>
#include <sys/stat.h>
#include <stdlib.h>

int main( void )
{
    int filedes, dup_filedes;

    filedes= open( "file",
                  O_WRONLY | O_CREAT | O_TRUNC,
                  S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP );

    if( filedes != -1 ) {
        dup_filedes = dup( filedes );
        if( dup_filedes != -1 ) {
            /* process file */
            /* ... */

            close( dup_filedes );
        }
        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*chsize(), close(), creat(), dup2(), eof(), errno, execl(), execle(),
execlp(), execlpe(), execv(), execve(), execvp(), execvpe(), fcntl(),
fileno(), fstat(), isatty(), lseek(), open(), read(), sopen(), stat(), tell(),
umask(), write()*

Synopsis:

```
#include <unistd.h>

int dup2( int filedes,
          int filedes2 );
```

Arguments:

filedes The file descriptor that you want to duplicate.

filedes The number that you want to use for the new file descriptor.

Library:

`libc`

Description:

The *dup2()* function duplicates the file descriptor specified by *filedes*. The number of the new file descriptor will be *filedes2*. If a file already is opened with this descriptor, the file is closed before the duplication is attempted.

The new file descriptor:

- references the same file or device
- has the same open mode (read and/or write)
- has an identical file position to the original (changing the position with one descriptor results in a changed position in the other).

Calling:

```
dup_filedes = dup2( filedes, filedes2 );
```

Is the same as:

```
close( filedes2 );
dup_filedes = fcntl( filedes , F_DUPFD, filedes2 );
```

Returns:

The value of *filedes2* for success, or -1 if an error occurs (*errno* is set).

Errors:

- EBADF The file descriptor, *filedes* isn't a valid open file descriptor, or *filedes2* is out of range.
- EMFILE There are already OPEN_MAX file descriptors in use.
- ENOSYS The *dup2()* function isn't implemented for the filesystem specified by *filedes*.

Examples:

```
#include <fcntl.h>
#include <unistd.h>
#include <sys/stat.h>
#include <stdlib.h>

int main( void )
{
    int filedes , dup_filedes ;

    filedes = open( "file",
        O_WRONLY | O_CREAT | O_TRUNC,
        S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP );

    if( filedes != -1 ) {
        dup_filedes = 4;
        if( dup2( filedes, dup_filedes ) != -1 ) {
            /* process file */
            /* ... */

            close( dup_filedes );
        }
        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*chsize(), close(), creat(), dup(), eof(), errno, execl(), execlp(),
execlpe(), execv(), execve(), execvp(), execvpe(), fcntl(), fileno(),
fstat(), isatty(), lseek(), open(), read(), sopen(), stat(), tell(), umask(),
write()*

eaccess()

© 2004, QNX Software Systems Ltd.

Check to see if a file or directory can be accessed (extended version)

Synopsis:

```
#include <libgen.h>
#include <unistd.h>

int eaccess( const char * path,
            int amode );
```

Arguments:

path The path to the file or directory that you want to access.

amode The access mode you want to check. This must be either:

- F_OK — test for file existence.

or a bitwise ORing of the following access permissions to be checked, as defined in the header `<unistd.h>`:

- R_OK — test for read permission.
- W_OK — test for write permission.
- X_OK — for a directory, test for search permission. Otherwise, test for execute permission.

Library:

`libc`

Description:

The `eaccess()` function is an extended version of `access()`. It checks if the file or directory specified by *path* exists and if it can be accessed with the file access permissions given by *amode*. However, unlike `access()`, it uses the effective user ID and effective group ID.

Returns:

0 The file or directory exists and can be accessed with the specified mode.

-1 An error occurred (*errno* is set.)

Errors:

EACCES	The permissions specified by <i>amode</i> are denied, or search permission is denied on a component of the path prefix.
EINVAL	An invalid value was specified for <i>amode</i> .
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of the <i>path</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENOENT	A component of the path isn't valid.
ENOSYS	The <i>eaccess()</i> function isn't implemented for the filesystem specified in <i>path</i> .
ENOTDIR	A component of the path prefix isn't a directory.
EROFS	Write access was requested for a file residing on a read-only file system.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

access(), chmod(), errno, fstat(), open(), stat()

Synopsis:

N/A

Description:

This linker symbol defines the end of the data segment, excluding BSS data. This variable isn't defined in any header file.

Classification:

QNX Neutrino

See also:

brk(), *_btext*, *_end*, *_etext*, *sbrk()*

encrypt()

© 2004, QNX Software Systems Ltd.

Encrypt a string

Synopsis:

```
#include <unistd.h>

void encrypt( char block[64],
             int flag );
```

Arguments:

block A 64-character array of binary values to encrypt. The function stores the encrypted value in the same array.

flag If the value of *flag* is 0, the function encrypts *block*; otherwise, *encrypt()* fails.

Library:

`libc`

Description:

The *encrypt()* function uses the NBS Data Encryption Standard (DES) algorithm and the key you specify by calling *setkey()* to encrypt the given block of data.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

crypt(), *setkey()*

end

© 2004, QNX Software Systems Ltd.

The end of the data segment, including BSS data

Synopsis:

N/A

Description:

This linker symbol defines the end of the data segment, including BSS data. This variable isn't defined in any header file.

Classification:

QNX Neutrino

See also:

brk(), *_btext*, *_edata*, *_etext*, *sbrk()*

Synopsis:

```
#include <grp.h>

int endgrent( void );
```

Library:

libc

Description:

The *endgrent()* routine closes the group name database file, so all group access routines behave as if *setgrent()* had never been called.

Returns:

Zero.

Classification:

Standard Unix, POSIX 1003.1g (draft)

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getgrent(), *setgrent()*

endhostent()

© 2004, QNX Software Systems Ltd.

Close the TCP connection and the hosts file

Synopsis:

```
#include <netdb.h>

void endhostent( void );
```

Library:

libsocket

Description:

The *endhostent()* routine closes the TCP connection and the hosts file.

Files:

/etc/hosts Host database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

gethostbyaddr(), *gethostbyname()*, *gethostent()*, **hostent**,
sethostent()

/etc/hosts, */etc/resolv.conf* in the *Utilities Reference*

Synopsis:

```
#include <gulliver.h>

uint16_t ENDIAN_BE16( uint16_t num );
```

Arguments:

num The big-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_BE16()* macro returns the native version of the big-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a big-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint16_t val = 0x1234;

    printf( "0x%04x = 0x%04x\n",
           val, ENDIAN_BE16( val ) );

    return EXIT_SUCCESS;
}
```

On a little-endian system, this prints:

0x1234 = 0x3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_BE16() is implemented as a macro.

See also:

ENDIAN_BE32(), *ENDIAN_BE64()*, *ENDIAN_LE16()*,
ENDIAN_LE32(), *ENDIAN_LE64()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint32_t ENDIAN_BE32( uint32_t num );
```

Arguments:

num The big-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_BE32()* macro returns the native version of the big-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a big-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint32_t val = 0xdeadbeef;

    printf( "0x%08x = 0x%08x\n",
           val, ENDIAN_BE32( val ) );

    return EXIT_SUCCESS;
}
```

On a little-endian system, this prints:

0xdeadbeef = 0xefbeadde

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_BE32() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE64()*, *ENDIAN_LE16()*,
ENDIAN_LE32(), *ENDIAN_LE64()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint64_t ENDIAN_BE64( uint64_t num );
```

Arguments:

num The big-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_BE64()* macro returns the native version of the big-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a big-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint64_t val = 0x1234deadbeef5678;

    printf( "0x%016Lx = 0x%016Lx\n",
           val, ENDIAN_BE64( val ) );

    return EXIT_SUCCESS;
}
```

On a little-endian system, this prints:

0x1234deadbeef5678 = 0x7856efbeadde3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_BE64() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_LE16()*,
ENDIAN_LE32(), *ENDIAN_LE64()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint16_t ENDIAN_LE16( uint16_t num );
```

Arguments:

num The little-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_LE16()* macro returns the native version of the little-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a little-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint16_t val = 0x1234;

    printf( "0x%04x = 0x%04x\n",
           val, ENDIAN_LE16( val ) );

    return EXIT_SUCCESS;
}
```

On a big-endian system, this prints:

0x1234 = 0x3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_LE16() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE32(), *ENDIAN_LE64()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint32_t ENDIAN_LE32( uint32_t num );
```

Arguments:

num The little-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_LE32()* macro returns the native version of the little-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a little-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint32_t val = 0xdeadbeef;

    printf( "0x%08x = 0x%08x\n",
           val, ENDIAN_LE32( val ) );

    return EXIT_SUCCESS;
}
```

On a big-endian system, this prints:

0xdeadbeef = 0xefbeadde

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_LE32() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE64()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint64_t ENDIAN_LE64( uint64_t num );
```

Arguments:

num The little-endian number you want to convert.

Library:

libc

Description:

The *ENDIAN_LE64()* macro returns the native version of the little-endian value *num*.

Returns:

The native-endian value of *num*.

Examples:

Convert a little-endian value to native-endian:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint64_t val = 0x1234deadbeef5678;

    printf( "0x%016Lx = 0x%016Lx\n",
           val, ENDIAN_LE64( val ) );

    return EXIT_SUCCESS;
}
```

On a big-endian system, this prints:

0x1234deadbeef5678 = 0x7856efbeadde3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_LE64() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_RET16()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint16_t ENDIAN_RET16( uint16_t num );
```

Arguments:

num The number you want to convert.

Library:

libc

Description:

The *ENDIAN_RET16()* macro returns the endian-swapped value of *num*.

Returns:

The endian-swapped value of *num*.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint16_t val = 0x1234;

    printf( "0x%04x = 0x%04x\n",
           val, ENDIAN_RET16( val ) );

    return EXIT_SUCCESS;
}
```

This prints:

0x1234 = 0x3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_RET16() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET32(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint32_t ENDIAN_RET32( uint32_t num );
```

Arguments:

num The number you want to convert.

Library:

libc

Description:

The *ENDIAN_RET32()* macro returns the endian-swapped value of *num*.

Returns:

The endian-swapped value of *num*.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint32_t val = 0xdeadbeef;

    printf( "0x%08x = 0x%08x\n",
           val, ENDIAN_RET32( val ) );

    return EXIT_SUCCESS;
}
```

This prints:

0xdeadbeef = 0xefbeadde

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_RET32() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET16(), *ENDIAN_RET64()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

uint64_t ENDIAN_RET64( uint64_t num );
```

Arguments:

num The number you want to convert.

Library:

libc

Description:

The *ENDIAN_RET64()* macro returns the endian-swapped value of *num*.

Returns:

The endian-swapped value of *num*.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint64_t val = 0x1234deadbeef5678;

    printf( "0x%016Lx = 0x%016Lx\n",
           val, ENDIAN_RET64( val ) );

    return EXIT_SUCCESS;
}
```

This prints:

0x1234deadbeef5678 = 0x7856efbeadde3412

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_RET64() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET16(), *ENDIAN_RET32()*, *ENDIAN_SWAP16()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

void ENDIAN_SWAP16( uint16_t * num );
```

Arguments:

num A pointer to the number you want to convert.

Library:

libc

Description:

The *ENDIAN_SWAP16()* macro endian-swaps the value pointed to by *num* in place.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint16_t val = 0x1234;
    ENDIAN_SWAP16( &val );

    printf( "val = 0x%04x\n", val );

    return EXIT_SUCCESS;
}
```

This prints:

```
val = 0x3412
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_SWAP16() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET16(), *ENDIAN_RET32()*, *ENDIAN_RET64()*,
ENDIAN_SWAP32(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

void ENDIAN_SWAP32( uint32_t * num );
```

Arguments:

num A pointer to the number you want to convert.

Library:

libc

Description:

The *ENDIAN_SWAP32()* macro endian-swaps the value pointed to by *num* in place.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint32_t val = 0xdeadbeef;
    ENDIAN_SWAP32( &val );

    printf( "val = 0x%08x\n", val );

    return EXIT_SUCCESS;
}
```

This prints:

```
val = 0xefbeadde
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_SWAP32() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET16(), *ENDIAN_RET32()*, *ENDIAN_RET64()*,
ENDIAN_SWAP16(), *ENDIAN_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_RET16()*, *UNALIGNED_RET32()*,
UNALIGNED_RET64(), *UNALIGNED_PUT16()*,
UNALIGNED_PUT32(), *UNALIGNED_PUT64()*

Synopsis:

```
#include <gulliver.h>

void ENDIAN_SWAP64( uint64_t * num );
```

Arguments:

num A pointer to the number you want to convert.

Library:

libc

Description:

The *ENDIAN_SWAP64()* macro endian-swaps the value pointed to by *num* in place.

Examples:

Swap the endianness of a value:

```
#include <stdio.h>
#include <stdlib.h>
#include <gulliver.h>
#include <inttypes.h>

int main( void )
{
    uint64_t val = 0x1234deadbeef5678LL;
    ENDIAN_SWAP16( &val );

    printf( "val = 0x%016x\n", val );

    return EXIT_SUCCESS;
}
```

This prints:

```
val = 0x7856efbeadde3412
```

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

Caveats:

ENDIAN_SWAP64() is implemented as a macro.

See also:

ENDIAN_BE16(), *ENDIAN_BE32()*, *ENDIAN_BE64()*,
ENDIAN_LE16(), *ENDIAN_LE32()*, *ENDIAN_LE64()*,
ENDIAN_RET16(), *ENDIAN_RET32()*, *ENDIAN_RET64()*,
ENDIAN_SWAP16(), *ENDIAN_SWAP32()*, *htonl()*, *htons()*, *ntohl()*,
ntohs(), *UNALIGNED_PUT16()*, *UNALIGNED_PUT32()*,
UNALIGNED_PUT64() *UNALIGNED_RET16()*,
UNALIGNED_RET32(), *UNALIGNED_RET64()*,

Synopsis:

```
#include <netdb.h>

void endnetent( void );
```

Library:

libsocket

Description:

The *endnetent()* routine closes the network name database file.

Files:

/etc/networks
Network name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getnetbyaddr(), *getnetbyname()*, *getnetent()*, **netent**, *setnetent()*
/etc/networks in the *Utilities Reference*

endprotoent()

© 2004, QNX Software Systems Ltd.

Close the protocol name database file

Synopsis:

```
#include <netdb.h>

void endprotoent( void );
```

Library:

libsocket

Description:

The *endprotoent()* routine closes the protocol name database file.

Files:

/etc/protocols
Protocol name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getprotobyname(), *getprotobynumber()*, *getprotoent()*, **protoent**, *setprotoent()*

/etc/protocols in the *Utilities Reference*

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

int endpwent( void );
```

Library:

libc

Description:

The *endpwent()* function closes the password name database file, so all password access routines behave as if *setpwent()* had never been called.

Returns:

Zero.

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getpwent(), *setpwent()*

endservent()

© 2004, QNX Software Systems Ltd.

Close the network services database file

Synopsis:

```
#include <netdb.h>

void endservent( void );
```

Library:

libsocket

Description:

The *endservent()* routine closes the network services database file.

Files:

/etc/services
Network services database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getservbyname(), *getservbyport()*, *getservent()*, **servent**, *setservent()*
/etc/services in the *Utilities Reference*

Synopsis:

```
#include <sys/types.h>
#include <shadow.h>

void endspent( void );
```

Library:

libc

Description:

The *endspent()* function closes the shadow password database file, so all password access routines behave as if *setspent()* had never been called.

Returns:

Zero.

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getspent(), setspent()

endutent()

© 2004, QNX Software Systems Ltd.

Close the current user-information file

Synopsis:

```
#include <utmp.h>

void endutent( void );
```

Library:

libc

Description:

The *endutent()* function closes the currently open file specified in `_PATH_UTMP`.

Files:

`_PATH_UTMP` The name of the user information file.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getutent(), *getutid()*, *getutline()*, *pututline()*, *setutent()*, `utmp`, *utmpname()*

`login` in the *Utilities Reference*

Synopsis:

```
#include <unistd.h>

extern char ** environ;
```

Library:

```
libc
```

Description:

When a process begins, an array of strings called the *environment* is made available. This array is pointed to by the external variable *environ*. The strings in the array have the form:

variable=value

and are terminated by a NULL pointer.

Classification:

POSIX 1003.1

Caveats:

Don't modify *environ* directly; use *clearenv()*, *getenv()*, *putenv()*, *searchenv()*, *setenv()*, and *unsetenv()*.

Changes to the environment affect all threads in a multithreaded process.

See also:

clearenv(), *getenv()*, *putenv()*, *searchenv()*, *setenv()*, *unsetenv()*

eof()

© 2004, QNX Software Systems Ltd.

Test if the end-of-file has been reached

Synopsis:

```
#include <unistd.h>

int eof( int filedes );
```

Arguments:

filedes A file descriptor for the file that you want to check.

Library:

`libc`

Description:

The *eof()* function is a low-level function that determines if the end of the file specified by *filedes* has been reached.

Input operations set the current file position; you can call the *eof()* function to detect the end of the file before more input operations to prevent attempts at reading beyond the end of the file.

Returns:

- 1 The end-of-file has been reached.
- 0 The end-of-file hasn't been reached.
- 1 An error occurred (*errno* is set).

Errors:

EBADF The file descriptor, *filedes*, isn't valid.

Examples:

```
#include <stdio.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>
```

```
int main( void )
{
    int filedes , len;
    char buffer[100];

    filedes = open( "file", O_RDONLY );
    if( filedes != -1 ) {
        while( ! eof( filedes ) ) {
            len = read( filedes , buffer, sizeof(buffer) - 1 );
            buffer[ len ] = '\0';
            printf( "%s", buffer );
        }
        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *feof()*, *open()*, *read()*

erand48()

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random **double** in a thread-safe manner*

Synopsis:

```
#include <stdlib.h>

double erand48( unsigned short int xsubi[3] );
```

Arguments:

xsubi An array that comprises the 48 bits of the initial value that you want to use.

Library:

libc

Description:

The *erand48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a nonnegative **double** uniformly distributed over the interval [0.0, 1.0]. It's a thread-safe version of *drand48()*.

The *xsubi* array should contain the desired initial value; this makes *erand48()* thread-safe, and lets you start a sequence of random numbers at any known value.

Returns:

A pseudo-random **double**.

Classification:

Standard Unix

Safety

Cancellation point No

Interrupt handler No

continued...

Safety

Signal handler	No
Thread	Yes

See also:

drand48(), jrand48(), lcong48(), lrand48(), mrand48(), nrand48(), seed48(), srand48()

erf()*, *erff()

© 2004, QNX Software Systems Ltd.

Compute the error function of a number

Synopsis:

```
#include <math.h>

double erf ( double x );

float erff ( float x );
```

Arguments:

x The number for which you want to compute the error function.

Library:

`libm`

Description:

The *erf()* and *erff()* functions compute the following:

$$\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

If *x* is large and the result of *erf()* is subtracted from 1.0, the results aren't very accurate; use *erfc()* instead.

This equality is true: $erf(-x) = -erf(x)$

Returns:

The value of the error function, or NAN if *x* is NAN.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Classification:

erf() is Standard Unix; *erff()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

erfc()

erfc()*, *erfcf()

© 2004, QNX Software Systems Ltd.

Complementary error function

Synopsis:

```
#include <math.h>

double erfc ( double x );

float erfcf ( float x );
```

Arguments:

x The number for which you want to compute the complementary error function.

Library:

libm

Description:

The *erfc()* and *erfcf()* functions calculate the complementary error function of *x* (i.e. the result of the error function, *erf()*, subtracted from 1.0). This is useful because the error function isn't very accurate when *x* is large.

The *erf()* function computes:

$$\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

This equality is true: $erfc(-x) = 2 - erf(x)$

Returns:

The value of the error function, or NAN if *x* is NAN.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Classification:

erfc() is standard Unix; *erfcf()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

erf()

err()*, *errx()

© 2004, QNX Software Systems Ltd.

Display a formatted error message, and then exit

Synopsis:

```
#include <err.h>

void err( int eval,
          const char *fmt, ...);

void errx( int eval,
           const char *fmt, ...);
```

Arguments:

eval The value to use as the exit code of the process.

fmt NULL, or a *printf()*-style string used to format the message.

Additional arguments
As required by the format string.

Library:

`libc`

Description:

The *err()* and *warn()* family of functions display a formatted error message on *stderr*:

- The functions without an **x** in their names display the string associated with the current value of *errno*.
- Those with a **v** are “varargs” functions.
- Those with **err** exit instead of returning.

Function	<i>errno</i> string?	Varargs?	Exits?
<i>err()</i>	Yes	No	Yes
<i>errx()</i>	No	No	Yes
<i>verr()</i>	Yes	Yes	Yes
<i>verrx()</i>	No	Yes	Yes
<i>vwarn()</i>	Yes	Yes	No
<i>vwarnx()</i>	No	Yes	No
<i>warn()</i>	Yes	No	No
<i>warnx()</i>	No	No	No

The *err()* function produces a message that consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, followed by a colon and a space, if the *fmt* argument isn't NULL
- the string associated with the current value of *errno*
- a newline character.

The *errx()* function produces a similar message, except that it doesn't include the string associated with *errno*. The message consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, if the *fmt* argument isn't NULL
- a newline character.

The *err()* and *errx()* functions don't return, but exit with the value of the argument *eval*.

Examples:

Display the current *errno* information string and exit:

```
if ((p = malloc(size)) == NULL)
    err(1, NULL);
if ((fd = open(file_name, O_RDONLY, 0)) == -1)
    err(1, "%s", file_name);
```

Display an error message and exit:

```
if (tm.tm_hour < START_TIME)
    errx(1, "too early, wait until %s", start_time_string);
```

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

printf(), *stderr*, *strerror()*, *verr()*, *verrx()*, *vwarn()*, *vwarnx()*, *warn()*, *warnx()*

Synopsis:

```
#include <errno.h>

extern int errno;

char * const sys_errlist[];
int sys_nerr;
```

Library:

libc

Description:

The *errno* variable is set to certain error values by many functions whenever an error has occurred.



You can't assume that the value of *errno* is valid unless the function that you've called indicates that an error has occurred. The runtime library never resets *errno* to 0.

The documentation for a function might list special meanings for certain values of *errno*, but this doesn't mean that these are the only values that the function might set.

The *errno* variable may be implemented as a macro, but you can always examine or set it as if it were a simple integer variable.



Each thread in a multi-threaded program has its own error value in its thread local storage. No matter which thread you're in, you can simply refer to *errno* — it's defined in such a way that it refers to the correct variable for the thread. For more information, see “Local storage for private data” in the documentation for *ThreadCreate()*.

The following variables are also defined in **<errno.h>**:

sys_errlist An array of error messages corresponding to *errno*.

sys_nerr The number of entries in the *sys_errlist* array.

The values for *errno* include at least the following values:

Value	Meaning
E2BIG	Argument list is too long
EACCES	Permission denied
EADDRINUSE	Address is already in use
EADDRNOTAVAIL	Can't assign requested address
EADV	Advertise error
EAFNOSUPPORT	Address family isn't supported by protocol family
EAGAIN	Resource is temporarily unavailable; try again
EALREADY	Operation is already in progress
EBADE	Invalid exchange
EBADF	Bad file descriptor
EBADFD	FD is invalid for this operation
EBADFSYS	Corrupted filesystem detected
EBADMSG	Bad message (1003.1b-1993)
EBADR	Invalid request descriptor
EBADRPC	RPC struct is bad
EBADRQC	Invalid request code
EBADSLT	Invalid slot
EBFONT	Bad font-file format
EBUSY	Device or resource is busy

continued...

Value	Meaning
ECANCELED	Operation canceled (1003.1b-1993)
ECHILD	No child processes
ECHRNG	Channel number is out of range
ECOMM	Communication error occurred on send
ECONNABORTED	Software caused connection to abort
ECONNREFUSED	Connection refused
ECONNRESET	Connection reset by peer
ECTRLTERM	Remap to the controlling terminal
EDEADLK	Resource deadlock avoided
EDEADLOCK	File locking deadlock
EDESTADDRREQ	Destination address is required
EDOM	Math argument is out of domain for the function
EDQUOT	Disk quota exceeded
EEXIST	File exists
EFAULT	Bad address
EFBIG	File is too large
EHOSTDOWN	Host is down
EHOSTUNREACH	Unable to communicate with remote node
EIDRM	Identifier removed
EILSEQ	Illegal byte sequence
EINPROGRESS	Operation now in progress
EINTR	Interrupted function call
EINVAL	Invalid argument

continued...

Value	Meaning
EIO	I/O error
EISCONN	Socket is already connected
EISDIR	Is a directory
EL2HLT	Level 2 halted
EL2NSYNC	Level 2 not synchronized
EL3HLT	Level 3 halted
EL3RST	Level 3 reset
ELIBACC	Can't access shared library
ELIBBAD	Accessing a corrupted shared library
ELIBEXEC	Attempting to exec a shared library
ELIBMAX	Attempting to link in too many libraries
ELIBSCN	The <code>.lib</code> section in <code>a.out</code> is corrupted
ELNRNG	Link number is out of range
ELOOP	Too many levels of symbolic links or prefixes
EMFILE	Too many open files
EMLINK	Too many links
EMORE	More to do, send message again
EMSGSIZE	Inappropriate message buffer length
EMULTIHOP	Multihop attempted
ENAMETOOLONG	Filename is too long
ENETDOWN	Network is down
ENETRESET	Network dropped connection on reset
ENETUNREACH	Network is unreachable

continued...

Value	Meaning
ENFILE	Too many open files in the system
ENOANO	No anode
ENOBUFFS	No buffer space available
ENOCSSI	No CSI structure available
ENODATA	No data (for no-delay I/O)
ENODEV	No such device
ENOENT	No such file or directory
ENOEXEC	Exec format error
ENOLCK	No locks available
ENOLIC	No license available
ENOLINK	The link has been severed
ENOMEM	Not enough memory
ENOMSG	No message of desired type
ENONDP	Need an NDP (8087...) to run
ENONET	Machine isn't on the network
ENOPKG	Package isn't installed
ENOPROTOPT	Protocol isn't available
ENOREMOTE	Must be done on local machine
ENOSPC	No space left on device
ENOSR	Out of streams resources
ENOSTR	Device isn't a stream
ENOSYS	Function isn't implemented
ENOTBLK	Block device is required

continued...

Value	Meaning
ENOTCONN	Socket isn't connected
ENOTDIR	Not a directory
ENOTEMPTY	Directory isn't empty
ENOTSOCK	Socket operation on nonsocket
ENOTSUP	Not supported (1003.1b-1993)
ENOTTY	Inappropriate I/O control operation
ENOTUNIQ	Given name isn't unique
ENXIO	No such device or address
EOK	No error
EOPNOTSUPP	Operation isn't supported
EOVERFLOW	Value too large to be stored in data type
EPERM	Operation isn't permitted
EPFNOSUPPORT	Protocol family isn't supported
EPIPE	Broken pipe
EPROCUNAVAIL	Bad procedure for program
EPROGMISMATCH	Program version wrong
EPROGUNAVAIL	RPC program isn't available
EPROTO	Protocol error
EPROTONOSUPPORT	Protocol isn't supported
EPROTOTYPE	Protocol is wrong type for socket
ERANGE	Result is too large
EREMCHG	Remote address changed
EREMOTE	The object is remote

continued...

Value	Meaning
ERESTART	Restartable system call
EROFS	Read-only filesystem
ERPCMISMATCH	RPC version is wrong
ESHUTDOWN	Can't send after socket shutdown
ESOCKTNOSUPPORT	Socket type isn't supported
ESPIPE	Illegal seek
ESRCH	No such process
ESRMNT	Server mount error
ESRVRFAULT	The receive side of a message transfer encountered a memory fault accessing the receive/reply buffer.
ESTALE	Potentially recoverable I/O error
ESTRPIPE	If pipe/FIFO, don't sleep in stream head
ETIME	Timer expired
ETIMEDOUT	Connection timed out
ETOOMANYREFS	Too many references: can't splice
ETXTBSY	Text file is busy
EUNATCH	Protocol driver isn't attached
EUSERS	Too many users (for UFS)
EWOULDBLOCK	Operation would block
EXDEV	Cross-device link
EXFULL	Exchange full

Examples:

```
/*
 * The following program makes an illegal call
 * to the write() function, then prints the
 * value held in errno.
 */
#include <stdio.h>
#include <unistd.h>
#include <errno.h>
#include <string.h>

int main( void )
{
    int errvalue;

    errno = EOK;
    write( -1, "hello, world\n",
          strlen( "hello, world\n" ) );
    errvalue = errno;
    printf( "The error generated was %d\n", errvalue );
    printf( "That means: %s\n", strerror( errvalue ) );
}
```

Classification:

POSIX 1003.1

See also:*h_errno, perror(), stderr, strerror()*

Synopsis:

Description:

This linker symbol defines the end of the text segment. This variable isn't defined in any header file.

Classification:

QNX Neutrino

See also:

brk(), *_btext*, *_edata*, *_end*, *sbrk()*

execl()

Execute a file

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

int execl( const char * path,
           const char * arg0,
           const char * arg1,
           :
           const char * argn,
           NULL );
```

Arguments:

<i>path</i>	The path of the file to execute.
<i>arg0, ..., argn</i>	Pointers to NULL-terminated character strings. These strings constitute the argument list available to the new process image. You must terminate the list with a NULL pointer. The <i>arg0</i> argument must point to a filename that's associated with the process being started.

Library:

libc

Description:

The *execl()* function replaces the current process image with a new process image specified by *path*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the *exec** functions are passed on to the new process image in the corresponding *main()* arguments.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *path* is on a filesystem mounted with the `ST_NOSUID` flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *path*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *path*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by *setuid()*).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID
- supplementary group IDs
- time left until an alarm clock signal (see *alarm()*)
- current working directory
- root directory
- file mode creation mask (see *umask()*)

- process signal mask (see *sigprocmask()*)
- pending signal (see *sigpending()*)
- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

If you call this function from a process with more than one thread, all of the threads are terminated and the new executable image is loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec** function failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execspe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes

continued...

Function	Description	POSIX?
<i>execve()</i>	NULL-terminated array of arguments, specify the new process environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process environment	No

Returns:

When *execl()* is successful, it doesn't return; otherwise, it returns -1 (*errno* is set).

Errors:

- E2BIG The argument list and the environment is larger than the system limit of ARG_MAX bytes.
- EACCESS The calling process doesn't have permission to search a directory listed in *path*, or it doesn't have permission to execute *path*, or *path*'s filesystem was mounted with the ST_NOEXEC flag.
- ELOOP Too many levels of symbolic links or prefixes.
- ENAMETOOLONG The length of *path* or an element of the **PATH** environment variable exceeds PATH_MAX.
- ENOENT One or more components of the pathname don't exist, or the *path* argument points to an empty string.
- ENOEXEC The new process image file has the correct access permissions, but isn't in the proper format.
- ENOMEM There's insufficient memory available to create the new process.

ENOTDIR A component of *path* isn't a directory.

Examples:

Replace the current process with **myprog** as if a user had typed:

```
myprog ARG1 ARG2
```

at the shell:

```
#include <stddef.h>
#include <process.h>

exec1( "myprog", "myprog", "ARG1", "ARG2", NULL );
```

In this example, **myprog** will be found if it exists in the current working directory.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

execle()

Execute a file

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

int execle( const char * path,
            const char * arg0,
            const char * arg1,
            :
            const char * argn,
            NULL,
            const char * envp [] );
```

Arguments:

<i>path</i>	The path of the file to execute.
<i>arg0, ..., argn</i>	Pointers to NULL-terminated character strings. These strings constitute the argument list available to the new process image. You must terminate the list with a NULL pointer. The <i>arg0</i> argument must point to a filename that's associated with the process being started.
<i>envp</i>	An array of character pointers to NULL-terminated strings. These strings constitute the environment for the new process image. Terminate the <i>envp</i> array with a NULL pointer.

Library:

libc

Description:

The *execle()* function replaces the current process image with a new process image specified by *path*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the exec functions are passed on to the new process image in the corresponding *main()* arguments.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling

process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *path* is on a filesystem mounted with the ST_NOSUID flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *path*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *path*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by *setuid()*).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID
- supplementary group IDs
- time left until an alarm clock signal (see *alarm()*)
- current working directory

- root directory
- file mode creation mask (see *umask()*)
- process signal mask (see *sigprocmask()*)
- pending signal (see *sigpending()*)
- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

A call to this function from a process with more than one thread results in all threads being terminated and the new executable image being loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec** function failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes

continued...

Function	Description	POSIX?
<i>execlpe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execle()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

- E2BIG The argument list and the environment is larger than the system limit of ARG_MAX bytes.
- EACCESS The calling process doesn't have permission to search a directory listed in *path*, or it doesn't have permission to execute *path*, or *path*'s filesystem was mounted with the ST_NOEXEC flag.
- ELOOP Too many levels of symbolic links or prefixes.
- ENAMETOOLONG The length of *path* or an element of the **PATH** environment variable exceeds PATH_MAX.
- ENOENT One or more components of the pathname don't exist, or the *path* argument points to an empty string.

ENOEXEC	The new process's image file has the correct access permissions, but isn't in the proper format.
ENOMEM	There's insufficient memory available to create the new process.
ENOTDIR	A component of <i>path</i> isn't a directory.

Examples:

Replace the current process with **myprog** as if a user had typed:

```
myprog ARG1 ARG2
```

at the shell:

```
#include <stddef.h>
#include <process.h>

char* env_list[] = { "SOURCE=MYDATA",
                    "TARGET=OUTPUT",
                    "lines=65",
                    NULL
                  };

execle( "myprog",
        "myprog", "ARG1", "ARG2", NULL,
        env_list );
```

In this example, **myprog** will be found if it exists in the current working directory. The environment for the invoked program consists of the three environment variables **SOURCE**, **TARGET** and **lines**.

Classification:

POSIX 1003.1

Safety

Cancellation point No

Interrupt handler No

continued...

Safety

Signal handler	Yes
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

Synopsis:

```
#include <process.h>

int execlp( const char * file,
            const char * arg0,
            const char * arg1,
            :
            const char * argn,
            NULL );
```

Arguments:

<i>file</i>	Used to construct a pathname that identifies the new process image file. If the <i>file</i> argument contains a slash character, the <i>file</i> argument is used as the pathname for the file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable PATH .
<i>arg0, ..., argn</i>	Pointers to NULL-terminated character strings. These strings constitute the argument list available to the new process image. Terminate the list terminated with a NULL pointer. The <i>arg0</i> argument must point to a filename that's associated with the process.

Library:

libc

Description:

The *execlp()* function replaces the current process image with a new process image specified by *file*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the exec functions are passed on to the new process image in the corresponding *main()* arguments.

If the process image file isn't a valid executable object, the contents of the file are passed as standard input to a command interpreter conforming to the *system()* function. In this case, the command interpreter becomes the new process image.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *file* is on a filesystem mounted with the ST_NOSUID flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *file*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *file*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by *setuid()*).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID

- supplementary group IDs
- time left until an alarm clock signal (see *alarm()*)
- current working directory
- root directory
- file mode creation mask (see *umask()*)
- process signal mask (see *sigprocmask()*)
- pending signal (see *sigpending()*)
- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

A call to this function from a process with more than one thread results in all threads being terminated and the new executable image being loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec** failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes

continued...

Function	Description	POSIX?
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execlpe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execlp()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

E2BIG	The argument list and the environment is larger than the system limit of ARG_MAX bytes.
EACCESS	The calling process doesn't have permission to search a directory listed in <i>file</i> , or it doesn't have permission to execute <i>file</i> , or <i>file</i> 's filesystem was mounted with the ST_NOEXEC flag.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of <i>file</i> or an element of the PATH environment variable exceeds PATH_MAX.

ENOENT	One or more components of the pathname don't exist, or the <i>file</i> argument points to an empty string.
ENOMEM	There's insufficient memory available to create the new process.
ENOTDIR	A component of <i>file</i> isn't a directory.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execle()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnvp()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

Synopsis:

```
#include <process.h>

int execlpe( const char * file,
             const char * arg0,
             const char * arg1,
             :
             const char * argn,
             NULL,
             const char * envp[] );
```

Arguments:

<i>file</i>	Used to construct a pathname that identifies the new process image file. If the <i>file</i> argument contains a slash character, the <i>file</i> argument is used as the pathname for the file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable PATH .
<i>arg0...argn</i>	Pointers to NULL-terminated character strings. These strings constitute the argument list available to the new process image. Terminate the list terminated with a NULL pointer. The <i>arg0</i> argument must point to a filename that's associated with the process.
<i>envp</i>	An array of character pointers to NULL-terminated strings. These strings constitute the environment for the new process image. Terminate the <i>envp</i> array with a NULL pointer.

Library:

libc

Description:



See *execl()* for further information on the *exec*()* family of functions.

The *execlpe()* function replaces the current process image with a new process image specified by *file*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.



If the new process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

The *execlpe()* function uses the paths listed in the **PATH** environment variable to locate the program to be loaded, provided that the following conditions are met:

- The argument *file* identifies the name of program to be loaded.
- If no path character (*/*) is included in the name, an attempt is made to load the program from one of the paths in the **PATH** environment variable.
- If **PATH** isn't defined, the current working directory is used.
- If a path character (*/*) is included in the name, the program is loaded from the path specified in *file*.

The process is started with the arguments specified in the NULL-terminated arguments *arg1...argn*. *arg0* should point to a filename associated with the program being loaded. Only *arg0* is required, *arg1...argn* are optional.

The new process's environment is specified in *envp*, a NULL-terminated array of NULL-terminated strings. *envp* cannot be NULL, but *envp[0]* can be a NULL pointer if no environment strings are passed.

Each pointer in *envp* points to a string in the form:

`variable=value`

that is used to define an environment variable.

The environment is the collection of environment variables whose values have been defined with the `export` shell command, the `env` utility, or by the successful execution of the `putenv()` or `setenv()` functions.

A program may read these values with the `getenv()` function.

An error is detected when the program cannot be found.

If the *file* is on a filesystem mounted with the `ST_NOSUID` flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process is set to the owner ID of *file*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *file*. The real user ID, real group ID and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process are saved as the saved set-user ID and the saved set-group ID used by `setuid()`.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execlpe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes

continued...

Function	Description	POSIX?
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execlpe()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

- E2BIG The argument list and the environment is larger than the system limit of ARG_MAX bytes.
- EACCESS The calling process doesn't have permission to search a directory listed in *file*, or it doesn't have permission to execute *file*, or *file*'s filesystem was mounted with the ST_NOEXEC flag.
- ELOOP Too many levels of symbolic links or prefixes.
- ENAMETOOLONG The length of *file* or an element of the **PATH** environment variable exceeds PATH_MAX.
- ENOENT One or more components of the pathname don't exist, or the *file* argument points to an empty string.
- ENOMEM There's insufficient memory available to create the new process.
- ENOTDIR A component of *file* isn't a directory.

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execle()*, *execlp()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

execv()

Execute a file

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

int execv( const char * path,
           char * const argv[] );
```

Arguments:

- path* A path name that identifies the new process image file.
- argv* An array of character pointers to NULL-terminated strings. Your application must ensure that the last member of this array is a NULL pointer. These strings constitute the argument list available to the new process image. The value in *argv[0]* must point to a filename that's associated with the process being started.

Library:

`libc`

Description:

The *execv()* function replaces the current process image with a new process image specified by *path*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the exec functions are passed on to the new process image in the corresponding *main()* arguments.

The environment for the new process image is taken from the external variable *environ* in the calling process.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *path* is on a filesystem mounted with the `ST_NOSUID` flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *path*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *path*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by `setuid()`).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID
- supplementary group IDs
- time left until an alarm clock signal (see `alarm()`)
- current working directory
- root directory
- file mode creation mask (see `umask()`)
- process signal mask (see `sigprocmask()`)
- pending signal (see `sigpending()`)

- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

A call to this function from a process with more than one thread results in all threads being terminated and the new executable image being loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec** failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execclpe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes

continued...

Function	Description	POSIX?
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execv()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

E2BIG	The argument list and the environment is larger than the system limit of ARG_MAX bytes.
EACCESS	The calling process doesn't have permission to search a directory listed in <i>path</i> , or it doesn't have permission to execute <i>path</i> , or <i>path</i> 's filesystem was mounted with the ST_NOEXEC flag.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of <i>path</i> or an element of the PATH environment variable exceeds PATH_MAX.
ENOENT	One or more components of the pathname don't exist, or the <i>path</i> argument points to an empty string.
ENOEXEC	The new process's image file has the correct access permissions, but isn't in the proper format.
ENOMEM	There's insufficient memory available to create the new process.
ENOTDIR	A component of <i>path</i> isn't a directory.

Examples:

```
#include <stddef.h>
#include <process.h>

char* arg_list[] = { "myprog", "ARG1", "ARG2", NULL };

execv( "myprog", arg_list );
```

The preceding invokes `myprog` as if the user entered:

```
myprog ARG1 ARG2
```

as a command at the shell. The program will be found if `myprog` exists in the current working directory.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execve()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

execve()

Execute a file

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

int execve( const char * path,
            char * const argv[],
            char * const envp[] );
```

Arguments:

- path* A path name that identifies the new process image file.
- argv* An array of character pointers to NULL-terminated strings. Your application must ensure that the last member of this array is a NULL pointer. These strings constitute the argument list available to the new process image. The value in *argv[0]* must point to a filename that's associated with the process being started.
- envp* An array of character pointers to NULL-terminated strings. These strings constitute the environment for the new process image. Terminate the *envp* array with a NULL pointer.

Library:

`libc`

Description:

The *execve()* function replaces the current process image with a new process image specified by *path*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```


where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the exec functions are passed on to the new process image in the corresponding *main()* arguments.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *path* is on a filesystem mounted with the `ST_NOSUID` flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *path*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *path*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by *setuid()*).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID
- supplementary group IDs
- time left until an alarm clock signal (see *alarm()*)
- current working directory
- root directory
- file mode creation mask (see *umask()*)

- process signal mask (see *sigprocmask()*)
- pending signal (see *sigpending()*)
- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

A call to this function from a process with more than one thread results in all threads being terminated and the new executable image being loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec** failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execspe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes

continued...

Function	Description	POSIX?
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execve()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

E2BIG	The argument list and the environment is larger than the system limit of ARG_MAX bytes.
EACCESS	The calling process doesn't have permission to search a directory listed in <i>path</i> , or it doesn't have permission to execute <i>path</i> , or <i>path</i> 's filesystem was mounted with the ST_NOEXEC flag.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of <i>path</i> or an element of the PATH environment variable exceeds PATH_MAX.
ENOENT	One or more components of the pathname don't exist, or the <i>path</i> argument points to an empty string.
ENOEXEC	The new process's image file has the correct access permissions, but isn't in the proper format.
ENOMEM	There's insufficient memory available to create the new process.

ENOTDIR A component of *path* isn't a directory.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execvp()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

execvp()

Execute a file

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

int execvp( const char * file,
            char * const argv[] );
```

Arguments:

- file* Used to construct a pathname that identifies the new process image file. If the *file* argument contains a slash character, the *file* argument is used as the pathname for the file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable **PATH**.
- argv* An array of character pointers to NULL-terminated strings. Your application must ensure that the last member of this array is a NULL pointer. These strings constitute the argument list available to the new process image. The value in *argv[0]* must point to a filename that's associated with the process being started.

Library:

libc

Description:

The *execvp()* function replaces the current process image with a new process image specified by *file*. The new image is constructed from a regular, executable file called the new process image file. No return is made because the calling process image is replaced by the new process image.

When a C-language program is executed as a result of this call, it's entered as a C-language function call as follows:

```
int main (int argc, char *argv[]);
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. In addition, the following variable:

```
extern char **environ;
```

is initialized as a pointer to an array of character pointers to the environment strings. The *argv* and *environ* arrays are each terminated by a null pointer. The null pointer terminating the *argv* array isn't counted in *argc*.

Multithreaded applications shouldn't use the *environ* variable to access or modify any environment variable while any other thread is concurrently modifying any environment variable. A call to any function dependent on any environment variable is considered a use of the *environ* variable to access that environment variable.

The arguments specified by a program with one of the exec functions are passed on to the new process image in the corresponding *main()* arguments.

If the process image file isn't a valid executable object, the contents of the file are passed as standard input to a command interpreter conforming to the *system()* function. In this case, the command interpreter becomes the new process image.

The environment for the new process image is taken from the external variable *environ* in the calling process.

The number of bytes available for the new process's combined argument and environment lists is ARG_MAX.

File descriptors open in the calling process image remain open in the new process image, except for when *fcntl()*'s FD_CLOEXEC flag is set. For those file descriptors that remain open, all attributes of the open file description, including file locks remain unchanged. If a file descriptor is closed for this reason, file locks are removed as described by *close()* while locks not affected by *close()* aren't changed.

Directory streams open in the calling process image are closed in the new process image.

Signals set to SIG_DFL in the calling process are set to the default action in the new process image. Signals set to SIG_IGN by the calling process images are ignored by the new process image. Signals set to be caught by the calling process image are set to the default action in the new process image. After a successful call, alternate signal stacks aren't preserved and the SA_ONSTACK flag is cleared for all signals.

After a successful call, any functions previously registered by *atexit()* are no longer registered.

If the *file* is on a filesystem mounted with the ST_NOSUID flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process image is set to the user ID of *file*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *file*. The real user ID, real group ID, and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process image are saved (as the saved set-user ID and the saved set-group ID used by *setuid()*).

Any shared memory segments attached to the calling process image aren't attached to the new process image.

The new process also inherits at least the following attributes from the calling process image:

- process ID
- parent process ID
- process group ID
- session membership
- real user ID
- real group ID
- supplementary group IDs

- time left until an alarm clock signal (see *alarm()*)
- current working directory
- root directory
- file mode creation mask (see *umask()*)
- process signal mask (see *sigprocmask()*)
- pending signal (see *sigpending()*)
- *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* (see *times()*)
- resource limits
- controlling terminal
- interval timers.

A call to this function from a process with more than one thread results in all threads being terminated and the new executable image being loaded and executed. No destructor functions are called.

Upon successful completion, the *st_atime* field of the file is marked for update. If the *exec* function failed but was able to locate the process image file, whether the *st_atime* field is marked for update is unspecified. On success, the process image file is considered to be opened with *open()*. The corresponding *close()* is considered to occur at a time after this open, but before process termination or successful completion of a subsequent call to one of the *exec** functions.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes

continued...

Function	Description	POSIX?
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execlpe()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execvp()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

E2BIG	The argument list and the environment is larger than the system limit of ARG_MAX bytes.
EACCESS	The calling process doesn't have permission to search a directory listed in <i>file</i> , or it doesn't have permission to execute <i>file</i> , or <i>file</i> 's filesystem was mounted with the ST_NOEXEC flag.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of <i>file</i> or an element of the PATH environment variable exceeds PATH_MAX.

ENOENT	One or more components of the pathname don't exist, or the <i>file</i> argument points to an empty string.
ENOMEM	There's insufficient memory available to create the new process.
ENOTDIR	A component of <i>file</i> isn't a directory.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), *atexit()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvpe()*, *_exit()*, *exit()*, *getenv()*, *main()*, *putenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*

execvpe()

© 2004, QNX Software Systems Ltd.

Execute a file

Synopsis:

```
#include <process.h>

int execvpe( const char * file,
             char * const argv[],
             char * const envp[] );
```

Arguments:

- file* Used to construct a pathname that identifies the new process image file. If the *file* argument contains a slash character, the *file* argument is used as the pathname for the file. Otherwise, the path prefix for this file is obtained by a search of the directories passed as the environment variable **PATH**.
- argv* An array of character pointers to NULL-terminated strings. Your application must ensure that the last member of this array is a NULL pointer. These strings constitute the argument list available to the new process image. The value in *argv[0]* must point to a filename that's associated with the process being started.
- envp* An array of character pointers to NULL-terminated strings. These strings constitute the environment for the new process image. Terminate the *envp* array with a NULL pointer.

Library:

libc

Description:



See *execl()* for further information on the *exec*()* family of functions.

The *execvpe()* function replaces the current process image with a new process image specified by *file*. The new image is constructed from a regular, executable file called the new process image file. No return is

made because the calling process image is replaced by the new process image.



If the new process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

The `execvpe()` function uses the paths listed in the **PATH** environment variable to locate the program to be loaded, provided that the following conditions are met:

- The *file* argument identifies the name of program to be loaded.
- If no path character (`/`) is included in the name, an attempt is made to load the program from one of the paths in the **PATH** environment variable.
- If **PATH** isn't defined, the current working directory is used.
- If a path character (`/`) is included in the name, the program is loaded from the path specified in *file*.

The process is started with the argument specified in *argv*, a NULL-terminated array of NULL-terminated strings. The *argv[0]* entry should point to a filename associated with the program being loaded. The *argv* argument can't be NULL but *argv[0]* can be NULL if no arguments are required.

The new process's environment is specified in *envp*, a NULL-terminated array of NULL-terminated strings. *envp* cannot be NULL, but *envp[0]* can be a NULL pointer if no environment strings are passed.

Each pointer in *envp* points to a string in the form:

```
variable=value
```

that is used to define an environment variable.

The environment is the collection of environment variables whose values have been defined with the **export** shell command, the **env**

utility, or by the successful execution of the *putenv()* or *setenv()* functions.

A program may read these values with the *getenv()* function.

An error is detected when the program cannot be found.

If the *file* is on a filesystem mounted with the *ST_NOSUID* flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the new process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the new process is set to the owner ID of *file*. Similarly, if the set-group ID mode bit is set, the effective group ID of the new process is set to the group ID of *file*. The real user ID, real group ID and supplementary group IDs of the new process remain the same as those of the calling process. The effective user ID and effective group ID of the new process are saved as the saved set-user ID and the saved set-group ID used by *setuid()*.

exec*() summary

Function	Description	POSIX?
<i>execl()</i>	NULL-terminated argument list	Yes
<i>execle()</i>	NULL-terminated argument list, specify the new process's environment	Yes
<i>execlp()</i>	NULL-terminated argument list, search for the new process in PATH	Yes
<i>execlope()</i>	NULL-terminated argument list, search for the new process in PATH , specify the new process's environment	No
<i>execv()</i>	NULL-terminated array of arguments	Yes
<i>execve()</i>	NULL-terminated array of arguments, specify the new process's environment	Yes
<i>execvp()</i>	NULL-terminated array of arguments, search for the new process in PATH	Yes

continued...

Function	Description	POSIX?
<i>execvpe()</i>	NULL-terminated array of arguments, search for the new process in PATH , specify the new process's environment	No

Returns:

When *execvpe()* is successful, it doesn't return; otherwise, it returns -1 and sets *errno*.

Errors:

E2BIG	The argument list and the environment is larger than the system limit of ARG_MAX bytes.
EACCESS	The calling process doesn't have permission to search a directory listed in <i>file</i> , or it doesn't have permission to execute <i>file</i> , or <i>file</i> 's filesystem was mounted with the ST_NOEXEC flag.
ELOOP	Too many levels of symbolic links or prefixes.
ENAMETOOLONG	The length of <i>file</i> or an element of the PATH environment variable exceeds PATH_MAX.
ENOENT	One or more components of the pathname don't exist, or the <i>file</i> argument points to an empty string.
ENOMEM	There's insufficient memory available to create the new process.
ENOTDIR	A component of <i>file</i> isn't a directory.

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*abort(), atexit(), errno, execl(), execl(), execlp(), execlpe(), execv(),
execve(), execvp(), _exit(), exit(), getenv(), main(), putenv(), spawn(),
spawnl(), spawnle(), spawnlp(), spawnlpe(), spawnp(), spawnv(),
spawnve(), spawnvp(), spawnvpe(), system()*

Synopsis:

```
#include <stdlib.h>

void _exit( int status );
```

Arguments:

status The exit status to use for the program.

Library:

`libc`

Description:

The `_exit()` function causes normal program termination to occur.



The functions registered with `atexit()` aren't called when you use `_exit()` to terminate a program. If you want those functions to be called, use `exit()` instead.

The `_exit()` function does the following when a process terminates for any reason:

- 1 Closes all open file descriptors and directory streams in the calling process.
- 2 Notifies the parent process of the calling process if the parent called `wait()` or `waitpid()`. The low-order 8 bits of *status* are made available to the parent via `wait()` or `waitpid()`.
- 3 Saves the exit *status* if the parent process of the calling process isn't executing a `wait()` or `waitpid()` function. If the parent calls `wait()` or `waitpid()` later, this status is returned immediately.
- 4 Sends a `SIGHUP` signal to the calling process's children; this can indirectly cause the children to exit if they don't handle `SIGHUP`. Children of a terminated process are assigned a new parent process.

- 5 Sends a SIGCHLD signal to the parent process.
- 6 Sends a SIGHUP signal to each process in the foreground process group if the calling process is the controlling process for the controlling terminal of that process group.
- 7 Disassociates the controlling terminal from the calling process's session if the process is a controlling process, allowing it to be acquired by a new controlling process.
- 8 If the process exiting causes a process group to become orphaned, and if any member of the newly-orphaned process group is stopped, then a SIGHUP signal followed by a SIGCONT signal is sent to each process in the newly-orphaned process group.

Returns:

The `_exit()` function doesn't return.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char *argv[] )
{
    FILE *fp;

    if( argc <= 1 ) {
        fprintf( stderr, "Missing argument\n" );
        exit( EXIT_FAILURE );
    }

    fp = fopen( argv[1], "r" );
    if( fp == NULL ) {
        fprintf( stderr, "Unable to open '%s'\n", argv[1] );
        _exit( EXIT_FAILURE );
    }
    fclose( fp );

    /*
     * At this point, calling _exit() is the same as calling
     * return EXIT_SUCCESS;...
     */
    _exit( EXIT_SUCCESS );
}
```

}

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*abort(), atexit(), close(), execl(), execlp(), execlpe(), execv(),
execve(), execvp(), execvpe(), exit(), getenv(), main(), putenv(),
sigaction(), signal(), spawn(), spawnl(), spawnle(), spawnlp(),
spawnlpe(), spawnp(), spawnv(), spawnve(), spawnvp(), spawnvpe(),
system(), wait(), waitpid()*

exit()

© 2004, QNX Software Systems Ltd.

Exit the calling program

Synopsis:

```
#include <stdlib.h>

void exit( int status );
```

Arguments:

status The exit status to use for the program.

Library:

`libc`

Description:

The *exit()* function causes the calling program to exit normally. When a program exits normally:

- 1 All functions registered with the *atexit()* function are called.
- 2 All open file streams (those opened by *fopen()*, *fdopen()*, *freopen()*, or *popen()*) are flushed and closed.
- 3 All temporary files created by the *tmpfile()* function are removed.
- 4 The return *status* is made available to the parent process; *status* is typically set to `EXIT_SUCCESS` to indicate successful termination and set to `EXIT_FAILURE` or some other value to indicate an error.

Returns:

The *exit()* function doesn't return.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( int argc, char *argv[] )
{
    FILE *fp;

    if( argc <= 1 ) {
        fprintf( stderr, "Missing argument\n" );
        exit( EXIT_FAILURE );
    }

    fp = fopen( argv[1], "r" );
    if( fp == NULL ) {
        fprintf( stderr, "Unable to open '%s'\n", argv[1] );
        exit( EXIT_FAILURE );
    }
    fclose( fp );
    exit( EXIT_SUCCESS );

    /*
     * You'll never get here; this prevents compiler
     * warnings about "function has no return value".
     */
    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abort(), atexit(), _exit(), main()

Synopsis:

```
#include <math.h>

double exp( double x );

float expf( float x );
```

Arguments:

x The number for which you want to calculate the exponential.

Library:

libm

Description:

The *exp()* function computes the exponential function of x (e^x).

A range error occurs if the magnitude of x is too large.

Returns:

The exponential value of x .



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", exp(.5) );

    return EXIT_SUCCESS;
}
```

produces the output:

1.648721

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

The value of *expm1(x)* may be more accurate than *exp(x) - 1.0* for small values of *x*.

See also:

errno, *expm1*, *log()*

Synopsis:

```
#include <math.h>

double expm1 ( double x );

float expm1f ( float x );
```

Arguments:

x The number for which you want to calculate the exponential minus one.

Library:

libm

Description:

The *expm1()* and *expm1f()* functions compute the exponential of *x*, minus 1 ($e^x - 1$).

A range error occurs if the magnitude of *x* is too large.

The value of *expm1(x)* may be more accurate than *exp(x) - 1.0* for small values of *x*.

The *expm1()* and *log1p()* functions are useful for financial calculations of $((1+x)^n - 1) / x$, namely:

```
expm1(n * log1p(x)) / x
```

when *x* is very small (for example, when performing calculations with a small daily interest rate). These functions also simplify writing accurate inverse hyperbolic functions.

Returns:

The exponential value of *x*, minus 1.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
    double a, b;

    a = 2;
    b = expm1(a);
    printf("(e ^ %f) -1 is %f \n", a, b);

    return(0);
}
```

produces the output:

```
(e ^ 2.000000) -1 is 6.389056
```

Classification:

expm1() is standard Unix; *expm1f()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

exp(), *log1p()*

fabs()*, *fabsf()

© 2004, QNX Software Systems Ltd.

Compute the absolute value of a double number

Synopsis:

```
#include <math.h>

double fabs( double x );

float fabsf( float x );
```

Arguments:

x The number you want the absolute value of.

Library:

`libm`

Description:

These functions compute the absolute value of *x*.

Returns:

The absolute value of *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f %f\n", fabs(.5), fabs(-.5) );
    return EXIT_SUCCESS;
}
```

produces the output:

0.500000 0.500000

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

abs(), *cabs()*, *labs()*

fcfgopen()

© 2004, QNX Software Systems Ltd.

Open a configuration file

Synopsis:

```
#include <cfgopen.h>

FILE * fcfgopen( const char * path,
                 const char * mode,
                 int location,
                 const char * historical,
                 char * namebuf,
                 int nblen );
```

Arguments:

<i>path</i>	The name of the configuration file that you want to open.
<i>mode</i>	A string that describes the mode to open in; see <i>fopen()</i> .
<i>location</i>	Flags that describe how the path is constructed. See “Search condition flags” in the documentation for <i>cfgopen()</i> .
<i>historical</i>	A optional file to open as a last resort if none of the criteria for finding the path is met. This string works like a path search order, and lets you search more than one location. You can also specify %H to substitute the hostname value into the string. Specify NULL to ignore this option.
<i>namebuf</i>	A buffer to save the pathname in. Specify NULL to ignore this option.
<i>nblen</i>	The length of the buffer pointed to by <i>namebuf</i> . Specify 0 to ignore this option.

Library:

libc

Description:

The *fcfgopen()* function is similar to *cfgopen()* with these exceptions:

- The `CFGFILE_NOFD` flag isn't valid.
- The values for *flags* described in *open()* aren't valid.

Returns:

A valid *fd* if `CFGFILE_NOFD` isn't specified, a nonnegative value if `CFGFILE_NOFD` is specified, or -1 if an error occurs.

Classification:

QNX Neutrino

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

cfgopen(), *confstr()*

`mib.txt`, `snmpd.conf` in the *Utilities Reference*

fchmod()

© 2004, QNX Software Systems Ltd.

Change the permissions for a file

Synopsis:

```
#include <sys/types.h>
#include <sys/stat.h>

int fchmod( int fd,
            mode_t mode );
```

Arguments:

fd A file descriptor for the file whose permissions you want to change.

mode The new permissions for the file. For more information, see “Access permissions” in the documentation for *stat()*.

Library:

`libc`

Description:

The *fchmod()* function changes the permissions for a file referred to by *fd* to be the settings in the mode given by *mode*.

If the effective user ID of the calling process is equal to the file owner, or the calling process has appropriate privileges (for example, the superuser), *fchmod()* sets the S_ISUID, S_ISGID and the file permission bits, defined in the `<sys/stat.h>` header file, from the corresponding bits in the *mode* argument. These bits define access permissions for the user associated with the file, the group associated with the file, and all others.

For a regular file, if the calling process doesn't have appropriate privileges, and if the group ID of the file doesn't match the effective group ID, the S_ISGID (set-group-ID on execution) bit in the file's mode is cleared upon successful return from *fchmod()*.

Changing the permissions has no any effect on any file descriptors for files that are already open.

If *fchmod()* succeeds, the `st_ctime` field of the file is marked for update.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EBADF Invalid file descriptor.
- ENOSYS The *fchmod()* function isn't implemented for the filesystem specified by *fd*.
- EPERM The effective user ID doesn't match the owner of the file, and the calling process doesn't have appropriate privileges.
- EROFS The referenced file resides on a read-only filesystem.

Examples:

```
/*
 * Change the permissions of a list of files
 * to be read/write by the owner only
 */
#include <stdio.h>
#include <fcntl.h>
#include <sys/types.h>
#include <sys/stat.h>

int main( int argc, char **argv )
{
    int i;
    int ecode = 0;
    int fd;

    for( i = 1; i < argc; i++ ) {
        if( ( fd = open( argv[i], O_RDONLY ) ) == -1 ) {
            perror( argv[i] );
            ecode++;
        }
        else if( fchmod( fd, S_IRUSR | S_IWUSR ) == -1 ) {
```

```
        perror( argv[i] );
        ecode++;
    }

    close( fd );
}
return ecode;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

chmod(), *chown()*, *errno*, *fchown()*, *fstat()*, *open()*, *stat()*

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

int fchown( int fd,
            uid_t owner,
            gid_t group );
```

Arguments:

fd A file descriptor for the file whose ownership you want to change.

owner The user ID of the new owner.

group The group ID of the new owner.

Library:

`libc`

Description:

The *fchown()* function changes the user ID and group ID of the file referenced by *fd* to be the numeric values contained in *owner* and *group*, respectively.

Only processes with an effective user ID equal to the user ID of the file, or with appropriate privileges (for example, the superuser) may change the ownership of a file.

The `_POSIX_CHOWN_RESTRICTED` flag is enforced. This means that only the superuser may change the ownership of a file. The group of a file may be changed by the superuser, or also by a process with the effective user ID equal to the user ID of the file, if (and only if) *owner* is equal to the user ID of the file and *group* is equal to the effective group ID of the calling process.

If the *fd* argument refers to a regular file, the set-user-ID (`S_ISUID`) and set-group-ID (`S_ISGID`) bits of the file mode are cleared if the function is successful.

If *fchown()* succeeds, the `st_ctime` field of the file is marked for update.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EBADF Invalid file descriptor.
- EPERM The effective user ID doesn't match the owner of the file, or the calling process doesn't have appropriate privileges.
- EROFS The named file resides on a read-only filesystem.

Examples:

```
/*
 * Change the ownership of a list of files
 * to the current user/group
 */
#include <stdio.h>
#include <fcntl.h>
#include <sys/types.h>
#include <unistd.h>

int main( int argc, char **argv )
{
    int i;
    int ecode = 0;
    int fd;

    for( i = 1; i < argc; i++ ) {
        if( ( fd = open( argv[i], O_RDONLY ) ) == -1 ) {
            perror( argv[i] );
            ecode++;
        }
        else if( fchown( fd, getuid(), getgid() ) == -1 ) {
            perror( argv[i] );
            ecode++;
        }
    }

    close( fd );
}
```

```
    }  
    return ecode;  
}
```

Classification:

POSIX 1003.1a

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

chmod(), *chown()*, *errno*, *fchmod()*, *fstat()*, *lchown()*, *open()*, *stat()*

fclose()

Close a stream

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <stdio.h>

int fclose( FILE* fp );
```

Arguments:

fp The stream you want to close.

Library:

libc

Description:

The *fclose()* function closes the stream specified by *fp*. Any unwritten, buffered data is flushed before the file is closed. Any unread, buffered data is discarded.

If the associated buffer was automatically allocated, it's deallocated.

Returns:

0 for success, or EOF if an error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;

    fp = fopen( "stdio.h", "r" );
    if( fp != NULL ) {
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fcloseall(), fdopen(), fopen(), freopen()

fcloseall()

© 2004, QNX Software Systems Ltd.

Close all open stream files

Synopsis:

```
#include <stdio.h>

int fcloseall( void );
```

Library:

libc

Description:

The *fcloseall()* function closes all open streams, except *stdin*, *stdout* and *stderr*. This includes streams created (and not yet closed) by *fdopen()*, *fopen()* and *freopen()*.

Returns:

0

Errors:

If an error occurs, *errno* is set.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    printf( "The number of files closed is %d\n", fcloseall() );
    return EXIT_SUCCESS;
}
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fclose(), fdopen(), fopen(), freopen()

fcntl()

© 2004, QNX Software Systems Ltd.

Provide control over an open file

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>
#include <fcntl.h>

int fcntl( int fdes,
           int cmd,
           ... );
```

Arguments:

fdes The descriptor for the file you want to control.

cmd The command to execute; see below.

Library:

libc

Description:

The *fcntl()* function provides control over the open file referenced by file descriptor *fdes*. To establish a lock with this function, open with write-only permission (O_WRONLY) or with read/write permission (O_RDWR).

The type of control is specified by the *cmd* argument, which may require a third data argument (*arg*). The *cmd* argument is defined in <fcntl.h>, and includes at least the following values:

FALLOCSP	If the size of the file is less than the number of bytes specified by the extra <i>arg</i> argument, extend the file with NUL characters.
FDUPFD	Allocate and return a new file descriptor that's the lowest numbered available (i.e. not already open) file descriptor greater than or equal to the third argument, <i>arg</i> , taken as an int . The new file

descriptor refers to the same file as *fildev*, and shares any locks.

F_FREESP Truncate the file to the size (in bytes) specified by the extra argument, *arg*.

F_GETFD Get the file descriptor flags associated with the file descriptor *fildev*. File descriptor flags are associated with a single file descriptor, and don't affect other file descriptors referring to the same file.

F_GETFL Get the file status flags and the file access modes associated with *fildev*. The flags and modes are defined in `<fcntl.h>`.

The file status flags (see *open()* for more detailed information) are:

- **O_APPEND** — Set append mode.
- **O_NONBLOCK** — No delay.

The file access modes are:

- **O_RDONLY** — Open for reading only.
- **O_RDWR** — Open for reading and writing.
- **O_WRONLY** — Open for writing only.

F_GETLK Get the first lock that blocks the lock description pointed to by the third argument, *arg*, taken as a pointer to type `struct flock` (defined in `<fcntl.h>`). For more information, see the `flock` structure section below. The information returned overwrites the information passed to *fcntl()* in the structure pointed to by *arg*.

If no lock is found that prevents this lock from being created, the structure is left unchanged, except for the lock type, which is set to **F_UNLCK**. If a lock is found, the *l_pid* member of the structure pointed to by *arg* is set to the process ID of the process holding the blocking lock and *l_whence* is set to **SEEK_SET**.

F_SETFD Set the file descriptor flags associated with *fdes* to the third argument, *arg*, taken as type **int**. See the above discussion for more details.

The only defined file descriptor flag is:

FD_CLOEXEC When this flag is clear, the file remains open across *spawn*()* or *exec*()* calls; else the file is closed.

F_SETFL Set the file status flags, as shown above, for the open file description associated with *fdes* from the corresponding bits in the third argument, *arg*, taken as type **int**. You can't use this function to change the file access mode. All bits set in *arg*, other than the file status bits, are ignored.

F_SETLK Set or clear a file segment lock, according to the lock description pointed to by the third argument, *arg*, taken as a pointer to type **struct flock**, as defined in the header file **<fcntl.h>**, and documented below. This command is used to create the following locks (defined in **<fcntl.h>**):

F_RDLCK Shared or read locks.

F_UNLCK Remove either type of lock.

F_WRLCK Exclusive or write locks.

If a lock can't be set, *fcntl()* returns immediately.

F_SETLKW This command is the same as **F_SETLK**, except that when a lock is blocked by other locks, the process waits until the request can be satisfied. If a signal that's to be caught is received while *fcntl()* is waiting for a region, the call is interrupted without performing the lock operation, and *fcntl()* returns -1 with *errno* set to **EINTR**.

flock structure

The `flock` structure contains at least the following members:

<code>short l_type</code>	One of F_RDLCK, F_WRLCK or F_UNLCK.
<code>short l_whence</code>	One of the following flags that specify where the relative offset, <code>l_start</code> , is measured from: SEEK_CUR Current seek position. SEEK_END End of file. SEEK_SET Start of file.
<code>off_t l_start</code>	Relative offset in bytes.
<code>off_t l_len</code>	Consecutive bytes to lock; if 0, then until EOF; if negative, the preceding bytes up to, but not including, the start byte.
<code>pid_t l_pid</code>	Process ID of the process holding the lock, returned when <code>cmd</code> is F_GETLK.

When a shared lock is set on a segment of a file, other processes can set shared locks on the same segment, or a portion of it. A shared lock prevents other processes from setting exclusive locks on any portion of the protected area. A request for a shared lock fails if the file was opened write-only.

An exclusive lock prevents any other process from setting a shared or an exclusive lock on a portion of the protected area. A request for an exclusive lock fails if the file was opened read-only.

Locks may start and extend beyond the current end of file, but may not start or extend before the beginning of the file; to attempt to do so is an error. A lock extends to “infinity” (the largest possible value for the file offset) if `l_len` is set to zero. If `l_whence` and `l_start` point to the beginning of the file, and `l_len` is zero, the entire file is locked.

The calling process may have only one type of lock set for each byte of a file. Before successfully returning from an `F_SETLK` or `F_SETLKW` request, the previous lock type (if any) for each byte in the specified lock region is replaced by the new lock type. All locks associated with a file for a given process are removed when a file descriptor for that file is closed by the process, or the process holding the file descriptor terminates. Locks aren't inherited by a child process using the `fork()` function. However, locks are inherited across `exec*()` or `spawn*()` calls.



A potential for deadlock occurs if a process controlling a locked region is put to sleep by attempting to lock another process's locked region. If the system detects that sleeping until a locked region is unlocked would cause a deadlock, `fcntl()` fails with `EDEADLK`. However, the system can't always detect deadlocks in the network case, and care should be exercised in the design of your application for this possibility.



Locking is a protocol designed for updating a file shared among concurrently running applications. Locks are only advisory, that is, they don't prevent an errant or poorly-designed application from overwriting a locked region of a shared file. An application should use locks to indicate regions of a file that are to be updated by the application, and it should respect the locks of other applications.

The following functions ignore locks:

- *chsize()*
 - *ltrunc()*
 - *open()*
 - *read()*
 - *sopen()*
 - *write()*
-

Returns:

-1 if an error occurred (*errno* is set). The successful return value(s) depend on the request type specified by *arg*, as shown in the following table:

F_DUPFD	A new file descriptor.
F_GETFD	Value of the file descriptor flags (never a negative value).
F_GETFL	Value of the file status flags and access modes as shown above (never a negative value).
F_GETLK	Value other than -1.
F_SETFD	Value other than -1.
F_SETFL	Value other than -1.

F_SETLK	Value other than -1.
F_SETLKW	Value other than -1.

Errors:

EAGAIN	The argument <i>cmd</i> is F_SETLK, the type of lock (<i>l_type</i>) is a shared lock (F_RDLCK), and the segment of a file to be locked is already exclusive-locked by another process, or the type is an exclusive lock and some portion of the segment of a file to be locked is already shared-locked or exclusive-locked by another process.
EBADF	The <i>fildev</i> argument isn't a valid file descriptor. The argument <i>cmd</i> is F_SETLK or F_SETLKW, the type of lock (<i>l_type</i>) is a shared lock (F_RDLCK), and <i>fildev</i> isn't a valid file descriptor open for reading. The argument <i>cmd</i> is F_SETLK or F_SETLKW, the type of lock (<i>l_type</i>) is an exclusive lock (F_WRLCK), and <i>fildev</i> isn't a valid file descriptor open for writing.
EDEADLK	The argument <i>cmd</i> is F_SETLKW, and a deadlock condition was detected.
EINTR	The argument <i>cmd</i> is F_SETLKW, and the function was interrupted by a signal.
EINVAL	The argument <i>cmd</i> is F_DUPFD, and the third argument is negative, or greater than the configured number of maximum open file descriptors per process. The argument <i>cmd</i> is F_GETLK, F_SETLK or F_SETLKW, and the data <i>arg</i> isn't valid, or <i>fildev</i> refers to a file that doesn't support locking.

EMFILE	The argument <i>cmd</i> is F_DUPFD, and the process has no unused file descriptors, or no file descriptors greater than or equal to <i>arg</i> are available.
ENOLCK	The argument <i>cmd</i> is F_SETLK or F_SETLKW, and satisfying the lock or unlock request causes the number of lock regions in the system to exceed the system-imposed limit.
E_OVERFLOW	One of the values to be returned can't be represented correctly.

Examples:

```
/*
 * This program makes "stdout" synchronous
 * to guarantee the data is recoverable
 * (if it's redirected to a file).
 */
#include <unistd.h>
#include <fcntl.h>
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    int flags, retval;

    flags = fcntl( STDOUT_FILENO, F_GETFL );

    flags |= O_DSYNC;

    retval = fcntl( STDOUT_FILENO, F_SETFL, flags );
    if( retval == -1 ) {
        printf( "error setting stdout flags\n" );
        return EXIT_FAILURE;
    }

    printf( "hello QNX world\n" );

    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Read the <i>Caveats</i>
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *fcntl()* function may be a cancellation point in the case of `F_DUPFD` (when dupping across the network), `F_GETFD`, and `F_SETFD`.

See also:

close(), *dup()*, *dup2()*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *open()*

Synopsis:

```
#include <unistd.h>

int fdatasync( int filedes );
```

Arguments:

filedes The descriptor of the file that you want to synchronize.

Library:

`libc`

Description:

The *fdatasync()* function forces all queued I/O operations for the file specified by the *filedes* file descriptor to finish, synchronizing the file's data.

This function is similar to *fsync()*, except that *fsync()* also guarantees the integrity of file information, such as access and modification times.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EBADF The specified *filedes* isn't a valid file descriptor open for writing.

EINVAL The implementation doesn't support synchronized I/O for the given file.

ENOSYS The *fdatasync()* function isn't supported for the filesystem specified by *filedes*.

Classification:

POSIX 1003.1 (Realtime Extensions)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

aio_fsync(), close(), fcntl(), fsync(), open(), read(), sync(), write()

Synopsis:

```
#include <stdio.h>

FILE* fdopen( int filedes,
              const char* mode );
```

Arguments:

filedes The file descriptor that you want to associate with a stream.

mode The mode specified when *filedes* was originally opened. For information, see *fopen()*, except modes beginning with **w** don't cause truncation of the file.

Library:

`libc`

Description:

The *fdopen()* function associates a stream with the file descriptor *filedes*, which represents an opened file or device.

The *filedes* argument is a file descriptor that was returned by one of *accept()*, *creat()*, *dup()*, *dup2()*, *fcntl()*, *open()*, *pipe()*, or *sopen()*.

The *fdopen()* function preserves the offset maximum previously set for the open file description corresponding to *filedes*.

Returns:

A file stream for success, or NULL if an error occurs (*errno* is set).

Errors:

EBADF The *filedes* argument is not a valid file descriptor.

EINVAL The *mode* argument is not a valid mode.

EMFILE Too many file descriptors are currently in use by this process.

ENOMEM There is no memory for FILE structure.

Examples:

```
#include <stdio.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main( void )
{
    int filedes ;
    FILE *fp;

    filedes = open( "file", O_RDONLY );
    if( filedes != -1 ) {
        fp = fdopen( filedes , "r" );
        if( fp != NULL ) {
            /* Also closes the underlying FD, filedes. */
            fclose( fp );
        }
    }
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

creat(), *dup()*, *dup2()*, *errno*, *fcntl()*, *fopen()*, *freopen()*, *open()*, *pipe()*,
sopen()

feof()

© 2004, QNX Software Systems Ltd.

Test a stream's end-of-file flag

Synopsis:

```
#include <stdio.h>

int feof( FILE* fp );
```

Arguments:

fp The stream you want to test.

Library:

libc

Description:

The *feof()* function tests the end-of-file flag for the stream specified by *fp*.

Because the end-of-file flag is set when an input operation attempts to read past the end-of-file, the *feof()* function detects the end-of-file only after an attempt is made to read beyond the end-of-file. Thus, if a file contains 10 lines, the *feof()* won't detect the end-of-file after the tenth line is read; it will detect the end-of-file on the next read operation.

Returns:

0 if the end-of-file flag isn't set, or nonzero if the end-of-file flag is set.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

void process_record( char *buf )
{
    printf( "%s\n", buf );
}

int main( void )
{
    FILE *fp;
    char buffer[100];
```



```
fp = fopen( "file", "r" );
fgets( buffer, sizeof( buffer ), fp );
while( ! feof( fp ) ) {
    process_record( buffer );
    fgets( buffer, sizeof( buffer ), fp );
}
fclose( fp );

return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

clearerr(), ferror(), fgetc(), fgetchar(), fgets(), fgetwc(), fgetws(), fopen(), freopen(), getc(), getc_unlocked(), getchar(), getchar_unlocked(), gets(), getw(), getwc(), getwchar(), perror(), read()

fclose()

© 2004, QNX Software Systems Ltd.

Test a stream's error flag

Synopsis:

```
#include <stdio.h>

int fclose( FILE* fp );
```

Arguments:

fp The stream whose error flag you want to test.

Library:

libc

Description:

The *fclose()* function tests the error flag for the stream specified by *fp*.

Returns:

0 if the error flag isn't set, or nonzero if the error flag is set.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        c = fgetc( fp );
        if( fclose( fp ) ) {
            printf( "Error reading file\n" );
        }
    }
    fclose( fp );

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

*clearerr(), feof(), fgetc(), fgetchar(), fgets(), fgetwc(), fgetws(), getc(),
getc_unlocked(), getchar(), getchar_unlocked(), gets(), getw(),
getwc(), getwchar(), perror(), strerror()*

fflush()

© 2004, QNX Software Systems Ltd.

Flush the buffers for a stream

Synopsis:

```
#include <stdio.h>

int fflush( FILE* fp );
```

Arguments:

fp NULL, or the stream whose buffers you want to flush.

Library:

libc

Description:

If the stream specified by *fp* is open for output or update, the *fflush()* function causes any buffered (see *setvbuf()*) but unwritten data to be written to the file descriptor associated with the stream (see *fileno()*).

If the file specified by *fp* is open for input or update, the *fflush()* function undoes the effect of any preceding **ungetc** operation on the stream.

If *fp* is NULL, all open streams are flushed.

Returns:

0 for success, or EOF if an error occurs (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    printf( "Press Enter to continue..." );
    fflush( stdout );
    getchar();

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fgetc(), fgets(), fileno(), flushall(), fopen(), getc(), gets(), setbuf(), setvbuf(), ungetc()

ffs()

© 2004, QNX Software Systems Ltd.

Find the first bit set in a bit string

Synopsis:

```
#include <strings.h>

int ffs( int value );
```

Arguments:

value The bit string.

Library:

`libc`

Description:

The *ffs()* function finds the first bit set in *value* and returns the index of that bit. Bits are numbered starting from 1, starting at the rightmost bit.

Returns:

The index of the first bit set, or 0 if *value* is zero.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Synopsis:

```
#include <stdio.h>

int fgetc( FILE* fp );
```

Arguments:

fp The stream from which you want to read a character.

Library:

libc

Description:

The *fgetc()* function reads the next character from the stream specified by *fp*.

Returns:

The next character from *fp*, cast as **(int) (unsigned char)**, or EOF if end-of-file has been reached or if an error occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        while( ( c = fgetc( fp ) ) != EOF ) {
            fputc( c, stdout );
        }
    }
}
```

```
        fclose( fp );  
        return EXIT_SUCCESS;  
    }  
    return EXIT_FAILURE;  
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, feof(), ferror(), fgetchar(), fgets(), fopen(), fputc(), getc(), gets(), ungetc()

Synopsis:

```
#include <stdio.h>

int fgetchar( void );
```

Library:

libc

Description:

The *fgetchar()* function is the same as *fgetc()* with an argument of *stdin*.

Returns:

The next character from *stdin*, cast as `(int) (unsigned char)`, EOF if end-of-file has been reached on *stdin* or if an error occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    fp = freopen( "file", "r", stdin );
    if( fp != NULL ) {
        while( (c = fgetchar()) != EOF ) {
            putchar(c);
        }
        fclose( fp );
    }

    return EXIT_SUCCESS;
}
```

```
        return EXIT_FAILURE;
    }
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, feof(), ferror(), fgetc(), fputchar(), getc(), getchar()

Synopsis:

```
#include <stdio.h>

int fgetpos( FILE* fp,
            fpos_t* pos );
```

Arguments:

fp The stream whose position you want to determine.

pos A pointer to a `fpos_t` object where the function can store the position.

Library:

```
libc
```

Description:

The `fgetpos()` function stores the current position of the stream *fp* in the `fpos_t` object specified by *pos*.

You can use the value stored in *pos* in a call to `fsetpos()` if you want to reposition the file to the position at the time of the `fgetpos()` call.

Returns:

0 for success, or nonzero if an error occurs (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    fpos_t position;
    char buffer[80];

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
```

```
        fgetpos( fp, &position ); /* get position */
        fgets( buffer, 80, fp ); /* read record */
        fsetpos( fp, &position ); /* set position */
        fgets( buffer, 80, fp ); /* read same record */
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fopen(), fseek(), fsetpos(), ftell()

Synopsis:

```
#include <stdio.h>

char* fgets( char* buf,
             size_t n,
             FILE* fp );
```

Arguments:

buf A pointer to a buffer in which *fgets()* can store the characters that it reads.

n The maximum number of characters to read.

fp The stream from which to read the characters.

Library:

libc

Description:

The *fgets()* function reads a string of characters from the stream specified by *fp*, and stores them in the array specified by *buf*.

It stops reading characters when:

- the end-of-file is reached
- Or:
- a newline (`'\n'`) character is read
- Or:
- *n*-1 characters have been read.

The newline character isn't discarded. A null character is placed immediately after the last character read into the array.



Don't assume that there's a newline character in every string that you read with *fgets()*. A newline character isn't present if there are more than *n-1* characters before the newline.

Also, a newline character might not appear as the last character in a file when the end-of-file is reached.

Returns:

The same pointer as *buf*, or NULL if the stream is at the end-of-file or an error occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    char buffer[80];

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        while( fgets( buffer, 80, fp ) != NULL ) {
            fputs( buffer, stdout );
        }
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, feof(), ferror(), fopen(), fputs(), getc(), gets(), fgetc()

fgetspent()

© 2004, QNX Software Systems Ltd.

Get an entry from the shadow password database

Synopsis:

```
#include <sys/types.h>
#include <shadow.h>

struct spwd* fgetspent( FILE* f );
```

Arguments:

f The stream from which to read the shadow password database.

Library:

`libc`

Description:

The *fgetspent()* works like the *getspent()* function but it assumes that it's reading from a file formatted like a shadow password database file. This function uses a static buffer that's overwritten by each call.



The *fgetspent()*, *getspent()*, and *getspnam()* functions share the same static buffer.

Returns:

A pointer to an object of type `struct spwd` containing the next entry from the password database. For more information about this structure, see *putspent()*.

Errors:

The *fgetspent()* function uses the following functions, and as a result *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*

- *fopen()*
- *fseek()*
- *rewind()*

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <pwd.h>
#include <shadow.h>

/*
 * This program loops, reading a entries from a file
 * (which is formatted in the shadow password way)
 * reading the next shadow password entry.
 * For example this_file /etc/shadow
 */

int main( int argc, char** argv )
{
    FILE*    fp;
    struct spwd* sp;

    if (argc < 2) {
        printf("%s filename \n", argv[0]);
        return(EXIT_FAILURE);
    }

    if (!(fp = fopen(argv[1], "r"))) {
        fprintf(stderr, "Can't open file %s \n", argv[1]);
        return(EXIT_FAILURE);
    }

    while( (sp = fgetspent(fp)) != (struct spwd *) 0 ) {
        printf( "Username: %s\n", sp->sp_namp );
        printf( "Password: %s\n", sp->sp_pwdp );
    }

    fclose(fp);
    return( EXIT_SUCCESS );
}
```

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

errno, *getgrent()*, *getlogin()*, *getpwnam()*, *getpwuid()*, *getspent()*,
getspnam(), *putspent()*

Synopsis:

```
#include <wchar.h>

wint_t fgetwc( FILE * fp );
```

Arguments:

fp The stream from which you want to read a character.

Library:

libc

Description:

The *fgetwc()* function reads the next wide character from the stream specified by *fp*.

Returns:

The next character from *fp*, cast as (**wint_t**) (**wchar_t**), or WEOF if end-of-file has been reached or if an error occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Errors:

EAGAIN	The O_NONBLOCK flag is set for <i>fp</i> and would have been blocked by this operation.
EBADF	The file descriptor for <i>fp</i> isn't valid for reading.
EINTR	A signal terminated the read operation; no data was transferred.
EIO	Either a physical I/O error has occurred, or the process is in the background and is being ignored or blocked.

EOVERFLOW Cannot read at or beyond the offset maximum for this stream.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *feof()*, *ferror()*, *fputwc()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter

Synopsis:

```
#include <wchar.h>

wchar_t * fgetws( wchar_t * buf,
                 int n,
                 FILE * fp );
```

Arguments:

- buf* A pointer to a buffer in which *fgetws()* can store the wide characters that it reads.
- n* The maximum number of characters to read.
- fp* The stream from which to read the characters.

Library:

`libc`

Description:

The *fgetws()* function reads a string of wide characters from the stream specified by *fp*, and stores them in the array specified by *buf*.

It stops reading wide characters when one of the following occurs:

- The end-of-file is reached.
- A newline (`'\n'`) character is read.
- *n*-1 characters have been read.

The *fgetws()* function places a NUL at the end of the string.



Don't assume all strings have newline characters. A newline character isn't present when more than $n-1$ characters occur before the newline.

Also, a newline character might not appear as the last character in a file when the end-of-file is reached.

Returns:

NULL Failure; the stream is at the end-of-file or an error occurred (*errno* is set).

buf Success.



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Errors:

EAGAIN The O_NONBLOCK flag is set for *fp* and would have been blocked by this operation.

EBADF The file descriptor for *fp* isn't valid for reading.

EINTR A signal terminated the read operation; no data was transferred.

EIO Either a physical I/O error has occurred, or the process is in the background and is being ignored or blocked.

E_OVERFLOW Cannot read at or beyond the offset maximum for this stream.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *feof()*, *ferror()*, *fputws()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

fileno()

© 2004, QNX Software Systems Ltd.

Return the file descriptor for a stream

Synopsis:

```
#include <stdio.h>

int fileno( FILE * stream );
```

Arguments:

stream The stream whose file descriptor you want to find.

Library:

`libc`

Description:

The *fileno()* function returns the file descriptor for the specified file *stream*. This file descriptor can be used in POSIX input/output calls anywhere the value returned by *open()* can be used.

To associate a stream with a file descriptor, call *fdopen()*.



In QNX Neutrino, the file descriptor is also the connection ID (*coid*) used by various Neutrino-specific functions.

The following symbolic values in `<unistd.h>` define the file descriptors associated with the C language *stdin*, *stdout*, and *stderr* streams:

STDIN_FILENO

Standard input file number, *stdin* (0)

STDOUT_FILENO

Standard output file number, *stdout* (1)

STDERR_FILENO

Standard error file number, *stderr* (2)

Returns:

A file descriptor, or -1 if an error occurs (*errno* is set).

Examples:

```
#include <stdlib.h>
#include <stdio.h>

int main( void )
{
    FILE *stream;

    stream = fopen( "file", "r" );
    if( stream != NULL ) {
        printf( "File number is %d.\n", fileno( stream ) );
        fclose( stream );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Produces output similar to:

```
File number is 7.
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, fdopen(), fopen(), open()

Synopsis:

```
#include <math.h>

int finite ( double x );

int finitf ( float x );
```

Arguments:

x The number you want to test.

Library:

libm

Description:

The *finite()* and *finitf()* functions determine if *x* is finite.

Returns:

True (1) The value of *x* is finite.
False (\neq 1) The value of *x* is infinity or NAN.

Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
    double a, b, c, d;

    a = 2;
    b = -0.5;
    c = NAN;
    fp_exception_mask(_FP_EXC_DIVZERO, 1);
    d = 1.0/0.0;
    printf("%f is %s \n", a, (finite(a)) ? "finite" : "not-finite");
```

```
printf("%f is %s \n", b, (finite(b)) ? "finite" : "not-finite");  
printf("%f is %s \n", c, (finite(c)) ? "finite" : "not-finite");  
printf("%f is %s \n", d, (finite(d)) ? "finite" : "not-finite");  
  
return(0);  
}
```

produces the output:

```
2.000000 is finite  
-0.500000 is finite  
NAN is not-finite  
Inf is not-finite
```

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

isnan()

Synopsis:

```
#include <unistd.h>

int flink( int fd,
           const char *path );
```

Arguments:

fd The file descriptor.

path The path you want to associate with the file descriptor.

Library:

libc

Description:

The *flink()* function assigns the pathname, *path*, to the file associated with the file descriptor, *fd*.

Returns:

0 Success.

-1 An error occurred; *errno* is set.

Errors:

EACCES A component of either path prefix denies search permission, or the link named by *path* is in a directory with a mode that denies write permission.

EBADF The file descriptor *fd* is invalid.

EEXIST The link named by *path* already exists.

ELOOP Too many levels of symbolic links or prefixes.

EMLINK	The number of links to the file would exceed LINK_MAX.
ENAMETOOLONG	The length of the <i>path</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENOENT	This error code can mean the following: <ul style="list-style-type: none">• A component of either path prefix doesn't exist.• The <i>path</i> points to an empty string.
ENOSPC	The directory that would contain the link can't be extended.
ENOSYS	The <i>flink()</i> function isn't implemented for the filesystem specified in <i>path</i> .
ENOTDIR	A component of either path prefix isn't a directory.
EROFS	The requested link requires writing in a directory on a read-only file system.
EXDEV	The link named by <i>path</i> is on a different logical disk.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

link()

flock()

© 2004, QNX Software Systems Ltd.

Apply or remove an advisory lock on an open file

Synopsis:

```
#include <fcntl.h>

int flock( int filedes,
           int operation );
```

Arguments:

filedes The file descriptor of an open file.

operation What you want to do to the file; see below.

Library:

`libc`

Description:

The *flock()* function applies or removes an advisory lock on the file associated with the open file descriptor *filedes*. To establish a lock with this function, open with write-only permission (O_WRONLY) or with read/write permission (O_RDWR).

A lock is applied by specifying one of the following values for *operation*:

LOCK_EX Exclusive lock.

LOCK_NB Don't block when locking. This may be ORed with LOCK_EX or LOCK_SH to give nonblocking behavior.

LOCK_SH Shared lock.

LOCK_UN Unlock an existing lock operation.

Advisory locks allow cooperating processes to perform consistent operations on files, but they don't guarantee consistency.

The locking mechanism allows two types of locks: *shared* and *exclusive*. At any time, multiple shared locks may be applied to a file,

but at no time are multiple exclusive, or both shared and exclusive, locks allowed simultaneously on a file.

A shared lock may be upgraded to an exclusive lock, and vice versa, by specifying the appropriate lock type. The previous lock is released and a new lock is applied (possibly after other processes have gained and released the lock).

Requesting a lock on an object that's already locked causes the caller to be blocked until the lock may be acquired. If you don't want the caller to be blocked, you can specify `LOCK_NB` in the operation to fail the call (*errno* is set to `EWOULDBLOCK`).



Locks are applied to files, not file descriptors. That is, file descriptors duplicated through *dup()* or *fork()* don't result in multiple instances of a lock, but rather multiple references to a single lock. If a process holding a lock on a file forks and the child explicitly unlocks the file, the parent loses its lock.

Processes blocked awaiting a lock may be awakened by signals.

Returns:

- 0 The operation was successful.
- 1 An error occurred (*errno* is set).

Errors:

- `EBADF` Invalid descriptor, *filedes*.
- `EINVAL` The argument operation doesn't include one of `LOCK_EX`, `LOCK_SH`, or `LOCK_UN`.
- `ENOMEM` The system can't allocate sufficient memory to store lock resources.
- `EOPNOTSUPP`
 - The *filedes* argument refers to an object other than a file.

EWOULDBLOCK

The file is locked and LOCK_NB was specified.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

fcntl(), lockf(), open()

Synopsis:

```
#include <stdio.h>

void flockfile( FILE* file );
```

Arguments:

file A pointer to the **FILE** object for the file you want to lock.

Library:

libc

Description:

The *flockfile()* function provides for explicit application-level locking of *stdio* (**FILE**) objects. This function can be used by a thread to delineate a sequence of I/O statements that are to be executed as a unit.

The *flockfile()* function is used by a thread to acquire ownership of a **FILE**.

The implementation acts as if there is a lock count associated with each **FILE**. This count is implicitly initialized to zero when the **FILE** is created. The **FILE** object is unlocked when the count is zero. When the count is positive, a single thread owns the **FILE**. When the *flockfile()* function is called, if the count is zero or if the count is positive and the caller owns the **FILE**, the count is incremented. Otherwise, the calling thread is suspended, waiting for the count to return to zero.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

*getc_unlocked(), getchar_unlocked(), putc_unlocked(),
putchar_unlocked()*

Synopsis:

```
#include <math.h>

double floor( double x );

float floorf( float x );
```

Arguments:

x The value you want to round.

Library:

`libm`

Description:

These functions compute the largest integer $\leq x$ (rounding towards the “floor”).

Returns:

The largest integer $\leq x$.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", floor( -3.14 ) );
    printf( "%f\n", floor( -3. ) );
    printf( "%f\n", floor( 0. ) );
}
```

```
printf( "%f\n", floor( 3.14 ) );  
printf( "%f\n", floor( 3. ) );  
  
return EXIT_SUCCESS;  
}
```

produces the output:

```
-4.000000  
-3.000000  
0.000000  
3.000000  
3.000000
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

ceil(), *fmod()*

Synopsis:

```
#include <stdio.h>

int flushall( void );
```

Library:

libc

Description:

The *flushall()* function flushes all buffers associated with open input/output streams. A subsequent read operation on an input stream reads new data from the associated stream.

Calling the *flushall()* function is equivalent to calling *fflush()* for all open streams.

Returns:

0 Success.
-1 An error occurred (*errno* is set).

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

The QNX 4 version of this function returns the number of streams flushed.

See also:

errno, fopen(), fflush()

Synopsis:

```
#include <math.h>

double fmod( double x,
             double y );

float fmodf( float x,
             float y );
```

Arguments:

x An arbitrary number.
y The modulus.

Library:

`libm`

Description:

The *fmod()* and *fmodf()* functions compute the floating-point residue of $x \pmod{y}$, which is the remainder of $x \div y$, even if the quotient $x \div y$ isn't representable.

Returns:

The residue, $x - (i \times y)$, for some integer i such that, if y is nonzero, the result has the same sign as x and a magnitude less than the magnitude of y .

If y is zero, the function returns 0.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main( void )
{
    printf( "%f\n", fmod( 4.5, 2.0 ) );
    printf( "%f\n", fmod( -4.5, 2.0 ) );
    printf( "%f\n", fmod( 4.5, -2.0 ) );
    printf( "%f\n", fmod( -4.5, -2.0 ) );

    return EXIT_SUCCESS;
}
```

produces the output:

```
0.500000
-0.500000
0.500000
-0.500000
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

ceil(), *div()*, *fabs()*, *floor()*

fnmatch()

© 2004, QNX Software Systems Ltd.

Check to see if a file or path name matches a pattern

Synopsis:

```
#include <fnmatch.h>

int fnmatch( const char* pat,
             const char* str,
             int flags );
```

Arguments:

pat The pattern to match; see “Pattern Matching Special Characters,” below.

str The string to match against the pattern.

flags Flags that modify interpretation of *pat* and *str*; a bitwise inclusive OR of these bits:

FNM_PATHNAME

If this is set, a slash character in *str* is explicitly matched by a slash in *pat*; it isn’t matched by either the asterisk or question mark special characters, or by a bracket expression.

FNM_PERIOD If this is set, a leading period in *str* matches a period in *pat*, where the definition of “leading” depends on FNM_PATHNAME:

- If FNM_PATHNAME is set, a period is leading if it’s the first character in *str*, or if it immediately follows a slash.
- If FNM_PATHNAME isn’t set, a period is leading only if it’s the first character in *str*.

FNM_QUOTE If this isn’t set, a backslash (\) in *pat* followed by another character matches that second character. If FNM_QUOTE is set, a backslash is treated as an ordinary character.

Library:

`libc`

Description:

The *fnmatch()* function checks the file or path name specified by the *str* argument to see if it matches the pattern specified by the *pat* argument.

Pattern Matching Special Characters

A pattern-matching special character that is quoted is a pattern that matches the special character itself. When not quoted, such special characters have special meaning in the specification of patterns. The pattern-matching special characters and the contexts in which they have their special meaning are as follows:

- ? Matches any printable or nonprintable collating element except <newline>.
- * Matches any string, including the null string.

[bracket_expr]

Matches a single collating element as per Regular Expression Bracket Expressions (1003.2 2.9.1.2) except that:

- The exclamation point character (!) replaces the circumflex character (^) in its role as a nonmatching list in the regular expression notation.
- The backslash is used as an escape character within bracket expressions.

The ?, * and [characters aren't special when used inside a bracket expression.

The concatenation of patterns matching a single character is a valid pattern that matches the concatenation of the single characters or collating elements matched by each of the concatenated patterns. For example, the pattern **a[bc]** matches the strings **ab** and **ac**.

The concatenation of one or more patterns matching a single character with one or more asterisks (*) is a valid pattern. In such patterns, each asterisk matches a string of zero or more characters, up to the first character that matches the character following the asterisk in the pattern. For example, the pattern **a*d** matches the strings **ad**, **abd**, and **abcd**, but not the string **abc**.

When an asterisk is the first or last character in a pattern, it matches zero or more characters that precede or follow the characters matched by the remainder of the pattern. For example, the pattern **a*d*** matches the strings **ad**, **abcd**, **abcdef**, **aaaad** and **adddd**; the pattern ***a*d** matches the strings **ad**, **abcd**, **efabcd**, **aaaad** and **adddd**.

Returns:

- 0 The *str* argument matches the pattern specified by *pat*.
- Nonzero The *str* argument doesn't match the pattern specified by *pat*.

Examples:

```
/*
 * The following example accepts a set of patterns
 * for filenames as argv[1..argc]. It reads lines
 * from standard input, and outputs the lines that
 * match any of the patterns.
 */
#include <stdio.h>
#include <fnmatch.h>
#include <stdlib.h>
#include <limits.h>

int main( int argc, char **argv )
{
    int i;
    char buffer[PATH_MAX+1];

    while( gets( buffer ) ) {
        for( i = 0; i < argc; i++ ) {
            if( fnmatch( argv[i], buffer, 0 ) == 0 ) {
                puts( buffer );
                break;
            }
        }
    }
}
```

```
    }  
    exit( EXIT_SUCCESS );  
}
```

Classification:

POSIX 1003.1a

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

regcomp()

fopen()*, *fopen64()

© 2004, QNX Software Systems Ltd.

Open a file stream

Synopsis:

```
#include <stdio.h>

FILE * fopen( const char * filename,
              const char * mode );

FILE * fopen64( const char * filename,
                const char * mode );
```

Arguments:

filename The name of the file that you want to open.

mode The access mode; see below.

Library:

libc

Description:

The *fopen()* and *fopen64()* functions open a file stream for the file specified by *filename*. The *mode* string begins with one of the following sequences:

- a** Append: create a new file or open the file for writing at its end.
- a+** Append: open the file or create it for update, writing at end-of-file; use the default file translation.
- r** Open the file for reading.
- r+** Open the file for update (reading and/or writing); use the default file translation.
- w** Create the file for writing, or truncate it to zero length.
- w+** Create the file for update, or truncate it to zero length; use the default file translation.

You can add the letter **b** to the end of any of the above sequences to indicate that the file is (or must be) a binary file (this is an ANSI requirement for portability to systems that make a distinction between text and binary files, such as DOS). Under QNX Neutrino, there's no difference between text files and binary files.

- Opening a file in read mode (**r** in the *mode*) fails if the file doesn't exist or can't be read.
- Opening a file in append mode (**a** in the *mode*) causes all subsequent writes to the file to be forced to the current end-of-file, regardless of previous calls to the *fseek()* function.
- When a file is opened with update mode (**+** in the *mode*), both input and output may be performed on the associated stream.



When using a stream in update mode, writing can't be followed by reading without an intervening call to *fflush()*, or to a file-positioning function (*fseek()*, *fsetpos()* or *rewind()*). Similarly, reading can't be followed by writing without an intervening call to a file-positioning function, unless the read resulted in end-of-file.

The largest value that can be represented correctly in an object of type `off_t` shall be established as the offset maximum in the open file description.

Returns:

A pointer to a file stream for success, or NULL if an error occurs (*errno* is set).

Errors:

EACCES	Search permission is denied on a component of the <i>filename</i> prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file doesn't exist and write permission is denied for the parent directory of the file to be created.
--------	--

EBADFSYS	While attempting to open the named file, either the file itself or a component of the <i>filename</i> prefix was found to be corrupted. A system failure — from which no automatic recovery is possible — occurred while the file was being written to, or while the directory was being updated. You’ll need to invoke appropriate systems-administration procedures to correct this situation before proceeding.
EBUSY	File access was denied due to a conflicting open (see <i>sopen()</i>).
EINTR	The <i>fopen()</i> operation was interrupted by a signal.
EINVAL	The value of the <i>mode</i> argument is not valid.
EISDIR	The named file is a directory, and the <i>mode</i> argument specifies write-only or read/write access.
ELOOP	Too many levels of symbolic links or prefixes.
EMFILE	Too many file descriptors are currently in use by this process.
ENAMETOOLONG	The length of the <i>filename</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENFILE	Too many files are currently open in the system.
ENOENT	Either the named file or the <i>filename</i> prefix doesn’t exist, or the <i>filename</i> argument points to an empty string.
ENOMEM	There is no memory for FILE structure.
ENOSPC	The directory or filesystem that would contain the new file can’t be extended.

ENOSYS	The <i>fopen()</i> function isn't implemented for the filesystem specified in <i>filename</i> .
ENOTDIR	A component of the <i>filename</i> prefix isn't a directory.
ENXIO	The media associated with the file has been removed (e.g. CD, floppy).
EOVERFLOW	The named file is a regular file and the size of the file can't be represented correctly in an object of type <code>off_t</code> .
EROFS	The named file resides on a read-only filesystem.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;

    fp = fopen( "report.dat", "r" );
    if( fp != NULL ) {
        /* rest of code goes here */
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

fopen() is ANSI, *fopen64()* is for large-file support

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fclose()*, *fcloseall()*, *fdopen()*, *freopen()*, *freopen64()*

Synopsis:

```
#include <sys/types.h>
#include <process.h>

pid_t fork( void );
```

Library:

libc

Description:

The *fork()* function creates a new process. The new process (child process) is an exact copy of the calling process (parent process), except for the following:

- The child process has a unique process ID.
- The child process has a different parent process ID (which is the process ID of the calling process).
- The child process has its own copy of the parent's file descriptors. Each of the child's file descriptors refers to the same open file description with the corresponding file descriptor of the parent.
- The child process has its own copy of the parent's open directory streams.
- The child process's values of *tms_utime*, *tms_stime*, *tms_cutime*, and *tms_cstime* are set to zero.
- File locks previously set by the parent aren't inherited by the child.
- Pending alarms are cleared for the child process.
- The set of signals pending for the child process is initialized to the empty set.

Returns:

A value of zero to the child process; and the process ID of the child process to the parent process. Both processes continue to execute from the *fork()* function. If an error occurs, *fork()* returns -1 to the parent and sets *errno*.

Errors:

EAGAIN	Insufficient resources are available to create the child process.
ENOMEM	The process requires more memory than the system is able to supply.
ENOSYS	The <i>fork()</i> function isn't implemented for this memory protection model. See also "Caveats," below.

Examples:

```
/*
 * This program executes the program and arguments
 * specified by argv[1..argc]. The standard input
 * of the executed program is converted to upper
 * case.
 */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <ctype.h>
#include <process.h>
#include <sys/wait.h>

int main( int argc, char **argv )
{
    pid_t pid;
    pid_t wpid;
    int fd[2];
    char buffer[80];
    int i, len;
    int status;

    if( pipe( fd ) == -1 ) {
        perror( "pipe" );
        return EXIT_FAILURE;
    }
}
```

```
if( ( pid = fork() ) == -1 ) {
    perror( "fork" );
    return EXIT_FAILURE;
}

if( pid == 0 ) {
    /* This is the child process.
     * Move read end of the pipe to stdin ( 0 ),
     * close any extraneous file descriptors,
     * then use exec to 'become' the command.
     */
    dup2( fd[0], 0 );
    close( fd[1] );
    execvp( argv[1], argv+1 );

/* This can only happen if exec fails; print message
 * and exit.
 */
    perror( argv[1] );
    return EXIT_FAILURE;
} else {
    /* This is the parent process.
     * Remove extraneous file descriptors,
     * read descriptor 0, write into pipe,
     * close pipe, and wait for child to die.
     */
    close( fd[0] );
    while( ( len = read( 0, buffer, sizeof( buffer ) )
        ) > 0 ) {
        for( i = 0; i < len; i++ ) {
            if( isupper( buffer[i] ) )
                buffer[i] = tolower( buffer[i] );
        }
        write( fd[1], buffer, len );
    }
    close( fd[1] );
    do {
        wpid = waitpid( pid, &status, 0 );
    } while( WIFEXITED( status ) == 0 );
    return WEXITSTATUS( status );
}
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Currently, *fork()* is supported only in single-threaded applications. If you create a thread and then call *fork()*, the function returns -1 and sets *errno* to ENOSYS.

See also:

errno, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *wait()*

Synopsis:

```
#include <unix.h>

pid_t forkpty( int *amaster,
               char *name,
               struct termios *termp,
               struct winsize *winp );
```

Arguments:

<i>amaster</i>	A pointer to a location where <i>forkpty()</i> can store the file descriptor of the master side of the pseudo-tty.
<i>name</i>	NULL, or a pointer to a buffer where <i>forkpty()</i> can store the filename of the slave side of the pseudo-tty.
<i>termp</i>	NULL, or a pointer to a termios structure that describes the terminal's control attributes to apply to the slave side of the pseudo-tty.
<i>winp</i>	A pointer to a winsize structure that defines the window size to use for the slave side of the pseudo-tty.

Library:

libc

Description:

The *forkpty()* function combines *openpty()*, *fork()*, and *login_tty()* to create a new process operating in a pseudo-tty.

This function fails if either *openpty()* or *fork()* fails.

Returns:

0 to the child process, the child's process ID to the parent, or -1 if an error occurred.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

fork(), *login_tty()*, *openpty()*, **termios**

Synopsis:

```
#include <fpstatus.h>

int fp_exception_mask ( int new_mask,
                       int set );
```

Arguments:

- new_mask* The new mask to apply. The bits include:
- `_FP_EXC_INVALID`
 - `_FP_EXC_DIVZERO`
 - `_FP_EXC_OVERFLOW`
 - `_FP_EXC_UNDERFLOW`
 - `_FP_EXC_INEXACT`
 - `_FP_EXC_DENORMAL`
- set* A value that indicates what you want the function to do:
- If *set* < 0, return the current mask. The *new_mask* argument is ignored.
 - If *set* = 0, disable the bits in the exception mask that correspond to the bits set in *new_mask*.
 - If *set* > 0, enable the bits in the exception mask that correspond to the bits set in *new_mask*.

Library:

`libm`

Description:

The *fp_exception_mask()* function gets or sets the current exception mask, depending on the value of the *set* argument.

Returns:

- If *set* < 0 The current exception mask.
- If *set* ≥ 0 The previous mask.



This function doesn't return a special value to indicate that an error occurred. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again.

Examples:

```
#include <fpstatus.h>

int main(int argc, char** argv)
{
    int ret;

    if ((ret = fp_exception_mask(0, -1)) < 0)
        printf("*** Problem retrieving exceptions \n");
    printf("Exceptions Enabled: \n\t");
    if (ret & _FP_EXC_INEXACT)
        printf("Inexact ");
    if (ret & _FP_EXC_DIVZERO)
        printf("DivZero ");
    if (ret & _FP_EXC_UNDERFLOW)
        printf("Underflow ");
    if (ret & _FP_EXC_OVERFLOW)
        printf("Overflow ");
    if (ret & _FP_EXC_INVALID)
        printf("Invalid ");
    printf("\n");

    /* Set the exception mask to enable division by zero errors */
    if ((ret = fp_exception_mask(_FP_EXC_DIVZERO, 1)) < 0)
        printf("*** Problem setting exceptions \n");
    if ((ret = fp_exception_mask(0, -1)) < 0)
        printf("*** Problem retrieving exceptions \n");
    printf("Exceptions Enabled: \n\t");
    if (ret & _FP_EXC_INEXACT)
        printf("Inexact ");
    if (ret & _FP_EXC_DIVZERO)
        printf("DivZero ");
    if (ret & _FP_EXC_UNDERFLOW)
        printf("Underflow ");
    if (ret & _FP_EXC_OVERFLOW)
```

```
    printf("Overflow ");
    if (ret & _FP_EXC_INVALID)
        printf("Invalid ");
    printf("\n");

    return(0);
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fp_exception_value(), *fp_precision()*, *fp_rounding()*

fp_exception_value()

© 2004, QNX Software Systems Ltd.

Get the value of the current exception registers

Synopsis:

```
#include <fpstatus.h>

int fp_exception_value( int mask );
```

Arguments:

mask A mask whose bits indicate which registers you want the value of. The bits include:

- `_FP_EXC_INVALID`
- `_FP_EXC_DIVZERO`
- `_FP_EXC_OVERFLOW`
- `_FP_EXC_UNDERFLOW`
- `_FP_EXC_INEXACT`
- `_FP_EXC_DENORMAL`

Library:

`libm`

Description:

The *fp_exception_value()* function gets the value of the current exception registers. Set bits indicate that the exception has signaled, unset bits indicate that the exception hasn't signaled.

Returns:

The value of the current exception registers based on the values from `<fpstatus.h>`.



This function doesn't return a special value to indicate that an error occurred. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again.

Examples:

```
#include <fpstatus.h>

int main(int argc, char** argv)
{
    int ret;

    /* Test to see if an operation has set (but not necessarily
     * signaled depending on the exception mask) the
     * division by zero bit:
     */

    if (fp_exception_value(_FP_EXC_DIVZERO) & _FP_EXC_DIVZERO)
        printf("Division by zero has occurred \n");
    else
        printf("Division by zero has not occurred \n");

    return(0);
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fp_precision(), *fp_rounding()*, *fp_exception_mask()*

Synopsis:

```
#include <fpstatus.h>

int fp_precision( int newprecision );
```

Arguments:

newprecision The new precision; one of:

- < 0 — return the current setting.
- `_FP_PREC_FLOAT`
- `_FP_PREC_DOUBLE`
- `_FP_PREC_EXTENDED`
- `_FP_PREC_DOUBLE_EXTENDED`

Library:

`libm`

Description:

The *fp_precision()* function sets or gets the current floating-point precision, depending on the value of *newprecision*.

Returns:

If *newprecision* is less than 0, the current precision; otherwise, the previous precision.



This function doesn't return a special value to indicate that an error occurred. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again.

Examples:

```
#include <fpstatus.h>

int main(int argc, char** argv)
{
    int ret;

    ret = fp_precision(-1);
    printf("Precision: ");
    if (ret == _FP_PREC_FLOAT)
        printf("Float \n");
    else if (ret == _FP_PREC_DOUBLE)
        printf("Double \n");
    else if (ret == _FP_PREC_EXTENDED)
        printf("Extended \n");
    else if (ret == _FP_PREC_DOUBLE_EXTENDED)
        printf("128 Bit \n");
    else if (ret == _FP_PREC_EXTENDED)
        printf("Extended \n");
    else if (ret == _FP_PREC_DOUBLE_EXTENDED)
        printf("128 Bit \n");
    else
        printf("Error \n");

    return(0);
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fp_exception_mask(), fp_exception_value(), fp_rounding()

fp_rounding()

© 2004, QNX Software Systems Ltd.

Set or get the current rounding

Synopsis:

```
#include <fpstatus.h>

int fp_rounding( int newrounding );
```

Arguments:

newrounding The new rounding; one of:

- < 0 — return the current setting.
- `_FP_ROUND_NEAREST`
- `_FP_ROUND_ZERO`
- `_FP_ROUND_POSITIVE`
- `_FP_ROUND_NEGATIVE`

Library:

`libm`

Description:

The *fp_rounding()* function sets or gets the current rounding mode, depending on the value of *newrounding*.

Returns:

If *newrounding* is less than 0, the current rounding mode; otherwise, the previous mode.



This function doesn't return a special value to indicate that an error occurred. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again.

Examples:

```
#include <fpstatus.h>
#include <stdlib.h>
#include <stdio.h>

int main(int argc, char** argv)
{
    int ret;

    ret = fp_rounding(-1);
    printf("Rounding mode: ");
    if (ret == _FP_ROUND_NEAREST)
        printf("Nearest \n");
    else if (ret == _FP_ROUND_POSITIVE)
        printf("Positive \n");
    else if (ret == _FP_ROUND_NEGATIVE)
        printf("Negative \n");
    else if (ret == _FP_ROUND_ZERO)
        printf("To Zero \n");
    else
        printf("Error \n");

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

fp_exception_mask(), fp_exception_value(), fp_precision()

Return the value of a configurable limit associated with a file

Synopsis:

```
#include <unistd.h>

long fpathconf( int filedes,
                int name );
```

Arguments:

filedes A file descriptor for the file whose limit you want to check.

name The name of the configurable limit; see below.

Library:

`libc`

Description:

The *fpathconf()* function returns a value of a configurable limit indicated by *name* that's associated with the file indicated by *filedes*.

Configurable limits are defined in `<confname.h>`, and include at least the following values:

`_PC_LINK_MAX`

Maximum value of a file's link count.

`_PC_MAX_CANON`

Maximum number of bytes in a terminal's canonical input buffer (edit buffer).

`_PC_MAX_INPUT`

Maximum number of bytes in a terminal's raw input buffer.

`_PC_NAME_MAX`

Maximum number of bytes in a file name (not including the terminating null).

`_PC_PATH_MAX`

Maximum number of bytes in a pathname (not including the terminating null).

`_PC_PIPE_BUF`

Maximum number of bytes that can be written atomically when writing to a pipe.

`_PC_CHOWN_RESTRICTED`

If defined (not -1), indicates that the use of the `chown()` function is restricted to a process with appropriate privileges, and to changing the group ID of a file to the effective group ID of the process or to one of its supplementary group IDs.

`_PC_NO_TRUNC`

If defined (not -1), indicates that the use of pathname components longer than the value given by `_PC_NAME_MAX` generates an error.

`_PC_VDISABLE`

If defined (not -1), this is the character value that can be used to individually disable special control characters in the `termios` control structure.

Returns:

The requested configurable limit, or -1 if an error occurred (*errno* is set).

Errors:

- `EINVAL` The *name* argument is invalid, or the indicated limit isn't supported for this *filedes*.
- `EBADF` The argument *filedes* is invalid.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main()
{
    long value;

    value = fpathconf( 0, _PC_MAX_INPUT );
    printf( "Input buffer size is %ld bytes\n",
           value );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

confstr(), *pathconf()*, *sysconf()*, **termios**

fprintf()

© 2004, QNX Software Systems Ltd.

Write output to a stream

Synopsis:

```
#include <stdio.h>

int fprintf( FILE* fp,
            const char* format,
            ... );
```

Arguments:

fp The stream to which you want to send the output.

format A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

Library:

libc

Description:

The *fprintf()* function writes output to the stream specified by *fp*, under control of the *format* specifier.

Returns:

The number of characters written, or a negative value if an output error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

char *weekday = { "Saturday" };
char *month = { "April" };

int main( void )
{
    fprintf( stdout, "%s, %s %d, %d\n",
            weekday, month, 10, 1999 );
}
```

```
        return EXIT_SUCCESS;
    }
```

Produces:

```
Saturday, April 10, 1999
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*, *vwprintf()*, *vfwprintf()*, *vprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*, *wprintf()*

fputc()

© 2004, QNX Software Systems Ltd.

Write a character to a stream

Synopsis:

```
#include <stdio.h>

int fputc( int c,
           FILE* fp );
```

Arguments:

c The character you want to write.

fp The stream you want to write the character to.

Library:

libc

Description:

The *fputc()* function writes the character specified by *c*, cast as **(int) (unsigned char)**, to the stream specified by *fp*.

Returns:

The character written, cast as **(int) (unsigned char)**, or EOF if an error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        while( ( c = fgetc( fp ) ) != EOF ) {
            fputc( c, stdout );
        }
        fclose( fp );
    }
```

```
        return EXIT_SUCCESS;
    }
    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

If *c* is negative, the value returned by this function isn't equal to *c* — unless *c* is -1 and an error occurred :-)

See also:

errno, *fgetc()*, *fopen()*, *fprintf()*, *fputchar()*, *fputs()*, *putc()*, *putchar()*, *puts()*

fputc()

© 2004, QNX Software Systems Ltd.

Write a character to stdout

Synopsis:

```
#include <stdio.h>

int fputc( int c );
```

Arguments:

c The character you want to write.

Library:

libc

Description:

The *fputc*() function writes the character specified by *c*, cast as (int) (unsigned char), to *stdout*. It's equivalent to *putc*() and to:

```
fputc( c, stdout );
```

Returns:

The character written, cast as (int) (unsigned char), or EOF if an error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        c = fgetc( fp );
        while( c != EOF ) {
```

```
        fputchar( c );
        c = fgetc( fp );
    }
    fclose( fp );

    return EXIT_SUCCESS;
}

return EXIT_FAILURE;
}
```

Classification:

QNX 4

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

If *c* is negative, the value returned by this function isn't equal to *c* — unless *c* is -1 and an error occurred :-)

See also:

errno, *fgetc()*, *fgetchar()*, *fprintf()*, *fputc()*, *fputs()*, *putc()*, *putchar()*

fputs()

© 2004, QNX Software Systems Ltd.

Write a string to an output stream

Synopsis:

```
#include <stdio.h>

int fputs( const char* buf,
           FILE* fp );
```

Arguments:

buf The string you want to write.

fp The stream you want to write the string to.

Library:

libc

Description:

The *fputs()* function writes the character string specified by *buf* to the output stream specified by *fp*.



The terminating NUL character isn't written.

Returns:

A nonnegative value for success, or EOF if an error occurs (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp_in, *fp_out;
    char buffer[80];

    fp_in = fopen( "file", "r" );
    fp_out = fopen( "outfile", "w" );
    if( fp_in != NULL && fp_out != NULL) {
```



```
        while( fgets( buffer, 80, fp_in ) != NULL ) {
            fputs( buffer, fp_out );
        }
        fclose( fp_in );
        fclose( fp_out );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fgets(), fopen(), fprintf(), fputc(), putc(), puts()

fputwc()

© 2004, QNX Software Systems Ltd.

Write a wide character to a stream

Synopsis:

```
#include <wchar.h>

wint_t fputwc( wchar_t wc,
               FILE * fp );
```

Arguments:

wc The wide character you want to write.

fp The stream you want to write the character to.

Library:

`libc`

Description:

The *fputwc()* function writes the wide character specified by *wc*, cast as `(wint_t) (wchar_t)`, to the stream specified by *fp*.

Returns:

The wide character written, cast as `(wint_t) (wchar_t)`, or WEOF if an error occurred (*errno* is set).



If *wc* exceeds the valid wide-character range, the value returned is the wide character written, not *wc*.

Errors:

EAGAIN The O_NONBLOCK flag is set for *fp* and would have been blocked by this operation.

EBADF The stream specified by *fp* isn't valid for writing.

EFBIG The file exceeds the maximum file size, the process's file size limit, or the function can't write at or beyond the offset maximum.

EINTR	A signal terminated the write operation; no data was transferred.
EIO	A physical I/O error has occurred or all of the following conditions were met: <ul style="list-style-type: none">• The process is in the background.• TOSTOP is set.• The process is blocking/ignoring SIGTTOU.• The process group is orphaned.
EPIPE	Can't write to pipe or FIFO because it's closed; a SIGPIPE signal is also sent to the thread.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:*errno, fgetwc(), fputws()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

fputws()

© 2004, QNX Software Systems Ltd.

Write a wide-character string to an output stream

Synopsis:

```
#include <wchar.h>

int fputws( const wchar_t * ws,
            FILE * fp );
```

Arguments:

buf The wide-character string you want to write.

fp The stream you want to write the string to.

Library:

`libc`

Description:

The *fputws()* function writes the wide-character string specified by *ws* to the output stream specified by *fp*.



The terminating NUL wide character isn't written.

Returns:

A nonnegative value for success, or WEOF if an error occurs (*errno* is set).

Errors:

EAGAIN The O_NONBLOCK flag is set for *fp* and would have been blocked by this operation.

EBADF The stream specified by *fp* isn't valid for writing.

EFBIG The file exceeds the maximum file size, the process's file size limit, or the function can't write at or beyond the offset maximum.

EINTR	A signal terminated the write operation; no data was transferred.
EIO	A physical I/O error has occurred or all of the following conditions were met: <ul style="list-style-type: none">• The process is in the background.• TOSTOP is set.• The process is blocking/ignoring SIGTTOU.• The process group is orphaned.
EPIPE	Can't write to pipe or FIFO because it's closed; a SIGPIPE signal is also sent to the thread.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:*errno, fgetws(), fputwc()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

fread()

© 2004, QNX Software Systems Ltd.

Read elements of a given size from a stream

Synopsis:

```
#include <stdio.h>

size_t fread( void* buf,
              size_t size,
              size_t num,
              FILE* fp );
```

Arguments:

buf A pointer to a buffer where the function can store the elements that it reads.

size The size of each element to read.

num The number of elements to read.

fp The stream from which to read the elements.

Library:

`libc`

Description:

The *fread()* function reads *num* elements of *size* bytes each from the stream specified by *fp* into the buffer specified by *buf*.

Returns:

The number of complete elements successfully read; this value may be less than the requested number of elements.

Use the *feof()* and *ferror()* functions to determine whether the end of the file was encountered or if an input/output error has occurred.

Errors:

If an error occurs, *errno* is set to indicate the type of error.

Examples:

The following example reads a simple student record containing binary data. The student record is described by the **struct student_data** declaration.

```
#include <stdio.h>
#include <stdlib.h>

struct student_data {
    int student_id;
    unsigned char marks[10];
};

size_t read_data( FILE *fp, struct student_data *p )
{
    return( fread( p, sizeof( struct student_data ), 1, fp ) );
}

int main( void )
{
    FILE *fp;
    struct student_data std;
    int i;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        while( read_data( fp, &std ) != 0 ) {
            printf( "id=%d ", std.student_id );

            for( i = 0; i < 10; i++ ) {
                printf( "%3d ", std.marks[ i ] );
            }

            printf( "\n" );
        }

        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fopen(), feof(), ferror()

Synopsis:

```
#include <stdlib.h>

void free( void* ptr );
```

Arguments:

ptr A pointer to the block of memory that you want to free. It's safe to call *free()* with a NULL pointer.

Library:

libc

Description:

The *free()* function deallocates the memory block specified by *ptr*, which was previously returned by a call to *calloc()*, *malloc()* or *realloc()*.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <malloc.h>

int main( void )
{
    char *buffer;

    buffer = (char *)malloc( 80 );
    if( buffer == NULL ) {
        printf( "Unable to allocate memory\n" );
        return EXIT_FAILURE;
    } else {
        /* rest of code goes here */

        free( buffer ); /* deallocate buffer */
    }

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

Calling *free()* on a pointer already deallocated by a call to *free()* or *realloc()* could corrupt the memory allocator's data structures.

See also:

alloca(), *calloc()*, *malloc()*, *realloc()*, *sbrk()*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netdb.h>

void freeaddrinfo( struct addrinfo * ai );
```

Arguments:

ai A pointer to the **addrinfo** structure that's at the beginning of the list to be freed.

Library:

libsocket

Description:

The *freeaddrinfo()* function frees the given list of **addrinfo** structures and the dynamic storage associated with each item in the list.

Classification:

POSIX 1003.1-2001

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

`addrinfo`, `gai_strerror()`, `getaddrinfo()`

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <ifaddr.h>

void freeifaddr( struct ifaddr * ifap );
```

Arguments:

ifap A pointer to the linked list of **ifaddr** structures to be freed.

Library:

libsocket

Description:

The *freeifaddr()* function frees the dynamically allocated data returned by *getifaddr()*.

Returns:

0 for success, or -1 if an error occurs (*errno* is set).

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *getifaddr()*, ***ifaddr***, *ioctl()*, *malloc()*, *socket()*, *sysctl()*

Synopsis:

```
#include <stdio.h>

FILE* freopen( const char* filename,
              const char* mode,
              FILE* fp );

FILE* freopen64( const char* filename,
                const char* mode,
                FILE* fp );
```

Arguments:

filename The name of the file to open.

mode The mode to use when opening the file. For more information, see *fopen()*.

fp The stream to associate with the file.

Library:

`libc`

Description:

The *freopen()* and *freopen64()* functions close the open stream *fp*, open the file specified by *filename*, and associate its stream with *fp*.

The largest value that can be represented correctly in an object of type `off_t` shall be established as the offset maximum in the open file description.

Returns:

A pointer to the newly opened stream, or NULL if an error occurs (*errno* is set).

Errors:

EACCES	Search permission is denied on a component of the <i>filename</i> prefix, or the file exists and the permissions specified by <i>mode</i> are denied, or the file doesn't exist and write permission is denied for the parent directory of the file to be created.
EBADFSYS	While attempting to open the named file, either the file itself or a component of the <i>filename</i> prefix was found to be corrupted. A system failure — from which no automatic recovery is possible — occurred while the file was being written to, or while the directory was being updated. You'll need to invoke appropriate systems-administration procedures to correct this situation before proceeding.
EBUSY	File access was denied due to a conflicting open (see <i>sopen()</i>).
EINTR	The <i>freopen()</i> operation was interrupted by a signal.
EINVAL	The value of the <i>mode</i> argument is not valid.
EISDIR	The named file is a directory, and the <i>mode</i> argument specifies write-only or read/write access.
ELOOP	Too many levels of symbolic links or prefixes.
EMFILE	Too many file descriptors are currently in use by this process.
ENAMETOOLONG	The length of the <i>filename</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.
ENFILE	Too many files are currently open in the system.

ENOENT	Either the named file or the <i>filename</i> prefix doesn't exist, or the <i>filename</i> argument points to an empty string.
ENOMEM	There is no memory for FILE structure.
ENOSPC	The directory or filesystem that would contain the new file can't be extended.
ENOSYS	The <i>freopen()</i> function isn't implemented for the filesystem specified in <i>filename</i> .
ENOTDIR	A component of the <i>filename</i> prefix isn't a directory.
ENXIO	The media associated with the file has been removed (e.g. CD, floppy).
EOVERFLOW	The named file is a regular file and the size of the file can't be represented correctly in an object of type <code>off_t</code> .
EROFS	The named file resides on a read-only filesystem.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE* fp;
    int c;

    /* Reopen the stdin stream so it's reading
     * from "file" instead of standard input.
     */
    fp = freopen( "file", "r", stdin );

    if( fp != NULL ) {
        /* Now we can read from "file" using the
         * stdin functions like fgetchar()...
         */
        while( ( c = fgetchar() ) != EOF ) {
            putchar( c );
        }
    }
}
```

```
    }  
    fclose( fp );  
    return EXIT_SUCCESS;  
}  
return EXIT_FAILURE;  
}
```

Classification:

freopen() is ANSI, *freopen64()* is for large-file support

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fclose()*, *fcloseall()*, *fdopen()*, *fopen()*, *fopen64()*

Break a floating-point number into a normalized fraction and an integral power of 2

Synopsis:

```
#include <math.h>

double frexp( double value,
              int* exp );

float frexpf( float value,
              int* exp );
```

Arguments:

value The value you want to break into a normalized fraction.

exp A pointer to a location where the function can store the integral power of 2.

Library:

`libm`

Description:

These functions break a floating-point number into a normalized fraction and an integral power of 2. It stores the integral power of 2 in the `int` pointed to by *exp*.

Returns:

x, such that *x* is a `double` with magnitude in the interval [0.5, 1] or 0, and *value* equals *x* times 2 raised to the power *exp*. If *value* is 0, then both parts of the result are 0.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main( void )
{
    int expon;
```

```
double value;

value = frexp( 4.25, &expon );
printf( "%f %d\n", value, expon );
value = frexp( -4.25, &expon );
printf( "%f %d\n", value, expon );

return EXIT_SUCCESS;
}
```

produces the output:

```
0.531250 3
-0.531250 3
```

Classification:

frexp() is ANSI; *frexpf()* is ANSI (draft)

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

ldexp(), *modf()*

Synopsis:

```
#include <stdio.h>

int fscanf( FILE* fp,
           const char* format,
           ... );
```

Arguments:

fp The stream that you want to read from.

format A string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

Library:

libc

Description:

The *fscanf()* function scans input from the stream specified by *fp*, under control of the argument *format*.

Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning reached the end of the input stream before storing any values (*errno* is set).

Examples:

Scan a date in the form “Friday March 26 1999”:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    int day;
```

```
int year;
char weekday[10];
char month[10];
FILE *in_data;

in_data = fopen( "file", "r" );
if( in_data != NULL ) {
    fscanf( in_data, "%s %s %d %d",
           weekday, month, &day, &year );

    printf( "Weekday=%s Month=%s Day=%d Year=%d\n",
           weekday, month, day, year );

    fclose( in_data );

    return EXIT_SUCCESS;
}
return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *vfscanf()*, *vfwscanf()*,
vscanf(), *vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

Synopsis:

```
#include <stdio.h>

int fseek( FILE* fp,
           long offset,
           int whence );

int fseeko( FILE* fp,
            off_t offset,
            int whence );
```

Arguments:

fp A **FILE** pointer returned by *fopen()* or *freopen()*.

offset The position to seek to, relative to one of the positions specified by *whence*.

whence The position from which to apply the offset; one of:

SEEK_SET	Compute the new file position relative to the start of the file. The value of <i>offset</i> must not be negative.
SEEK_CUR	Compute the new file position relative to the current file position. The value of <i>offset</i> may be positive, negative or zero. A SEEK_CUR with a 0 <i>offset</i> is necessary when you want to switch from reading to writing on a stream opened for updates.
SEEK_END	Compute the new file position relative to the end of the file.

Library:

libc

Description:

The *fseek()* function changes the current position of the stream specified by *fp*. This position defines the character that will be read or written by the next I/O operation on the file.

The *fseek()* function clears the end-of-file indicator, and undoes any effects of the *ungetc()* function on the stream.

You can use *ftell()* to get the current position of the stream before changing it. You can restore the position by using the value returned by *ftell()* in a subsequent call to *fseek()* with the *whence* parameter set to *SEEK_SET*.

Returns:

0 for success, or nonzero if an error occurs.

Errors:

If an error occurs, *errno* is set to indicate the type of error.

Examples:

Determine the size of a file, by saving and restoring the current position of the file:

```
#include <stdio.h>
#include <stdlib.h>

long filesize( FILE *fp )
{
    long int save_pos;
    long size_of_file;

    /* Save the current position. */
    save_pos = ftell( fp );

    /* Jump to the end of the file. */
    fseek( fp, 0L, SEEK_END );

    /* Get the end position. */
    size_of_file = ftell( fp );

    /* Jump back to the original position. */
```



```
        fseek( fp, save_pos, SEEK_SET );

        return( size_of_file );
    }

int main( void )
{
    FILE *fp;

    fp = fopen( "file", "r" );

    if( fp != NULL ) {
        printf( "File size=%ld\n", filesize( fp ) );
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

fseek() is ANSI; *fseeko()* is standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, *fgetpos()*, *fopen()*, *fsetpos()*, *ftell()*

fsetpos()

© 2004, QNX Software Systems Ltd.

Set the current position of a file

Synopsis:

```
#include <stdio.h>

int fsetpos( FILE* fp,
            const fpos_t* pos );
```

Arguments:

fp The stream whose position you want to set.

pos A pointer to a `fpos_t` object that specifies the new position for the stream. You must have initialized the value pointed to by *pos* by calling `fgetpos()` on the same file.

Library:

`libc`

Description:

The `fsetpos()` function sets the current position of the stream specified by *fp* according to the value of the `fpos_t` object pointed to by *pos*.

Returns:

0 for success, or nonzero if an error occurs (*errno* is set).

Examples:

See `fgetpos()`.

Classification:

ANSI

Safety

Cancellation point No

continued...

Safety

Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, fgetpos(), fopen(), fseek(), ftell()

fstat()*, *fstat64()

© 2004, QNX Software Systems Ltd.

Get file information, given a file description

Synopsis:

```
#include <sys/types.h>
#include <sys/stat.h>

int fstat( int filedes,
           struct stat* buf );

int fstat64( int filedes,
             struct stat64* buf );
```

Arguments:

filedes The descriptor of the file that you want to get information about.

buf A pointer to a buffer where the function can store the information about the file.

Library:

`libc`

Description:

The *fstat()* and *fstat64()* functions get information from the file specified by *filedes* and stores it in the structure pointed to by *buf*.

The file `<sys/stat.h>` contains definitions for `struct stat`, as well as following macros:

S_ISBLK(m) Test for block special file.

S_ISCHR(m) Test for character special file.

S_ISDIR(m) Test for directory.

S_ISFIFO(m) Test for FIFO.

S_ISLNK(m) Test for symbolic link.

- S_ISREG(m)* Test for regular file.
- S_TYPEISMQ(buf)*
 Test for message queue.
- S_TYPEISSEM(buf)*
 Test for semaphore.
- S_TYPEISSHM(buf)*
 Test for shared memory object.

The arguments to the macros are:

- m* The value of *st_mode* in a **stat** structure.
- buf* A pointer to a **stat** structure.

The macros evaluate to nonzero if the test is true, and zero if the test is false.

Access permissions are specified as a combination of bits in the *st_mode* field of the **stat** structure. These bits are defined in `<sys/stat.h>`. For more information, see “Access permissions” in the documentation for *stat()*.

The *st_mode* field also encodes the following bits:

- S_ISUID** Set user ID on execution. The process’s effective user ID (EUID) is set to that of the owner of the file when the file is run as a program. On a regular file, this bit may be cleared for security reasons on any write.
- S_ISGID** Set group ID on execution. Set effective group ID (EGID) on the process to the file’s group when the file is run as a program. On a regular file, this bit may be cleared for security reasons on any write.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EBADF The *filedes* argument isn't a valid file descriptor.
- ENOSYS The *fstat()* function isn't implemented for the filesystem specified by *filedes*.
- EOVERFLOW The file size in bytes or the number of blocks allocated to the file or the file serial number can't be represented correctly in the structure pointed to by *buf*.

Examples:

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main( void )
{
    int filedes;
    int rc;
    struct stat buf;

    filedes = open( "file", O_RDONLY );
    if( filedes != -1 ) {
        rc = fstat( filedes , &buf );
        if( rc != -1 ) {
            printf( "File size = %d\n", buf.st_size );
        }

        close( filedes );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

fstat() is POSIX 1003.1; *fstat64()* is for large-file support

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

creat(), *dup()*, *dup2()*, *errno*, *fcntl()*, *lstat()*, *open()*, *pipe()*, *sopen()*, *stat()*

fstatvfs()*, *fstatvfs64()

© 2004, QNX Software Systems Ltd.

Get filesystem information, given a file descriptor

Synopsis:

```
#include <sys/statvfs.h>

int fstatvfs( int fildev,
              struct statvfs *buf );

int fstatvfs64( int fildev,
                struct statvfs64 *buf );
```

Arguments:

fildev The descriptor for a file that resides on the filesystem that you want to get information about.

buf A pointer to a buffer where the function can store information about the filesystem; see below.

Library:

`libc`

Description:

The *fstatvfs()* function returns a “generic superblock” describing a filesystem; you can use it to get information about mounted filesystems. The *fstatvfs64()* function is a 64-bit version of *fstatvfs()*.

The *fildev* argument is an open file descriptor, obtained from a successful call to *open()*, *creat()*, *dup()*, *fcntl()*, or *pipe()*, for a file that resides on that filesystem. The filesystem type is known to the operating system. Read, write, or execute permission for the named file isn't required.

The *buf* argument is a pointer to a `statvfs` or `statvfs64` structure that's filled by the function. It contains at least:

```
unsigned long f_bsize
```

The preferred filesystem blocksize.

unsigned long *f_frsize*

The fundamental filesystem blocksize (if supported)

fsblkcnt_t *f_blocks*

The total number of blocks on the filesystem, in units of *f_frsize*.

fsblkcnt_t *f_bfree*

The total number of free blocks.

fsblkcnt_t *f_bavail*

The number of free blocks available to a nonsuperuser.

fsfilcnt_t *f_files*

The total number of file nodes (inodes).

fsfilcnt_t *f_ffree*

The total number of free file nodes.

fsfilcnt_t *f_favail*

The number of inodes available to a nonsuperuser.

unsigned long *f_fsid*

The filesystem ID (dev for now).

char *f_basetype*[16]

The type of the target filesystem, as a null-terminated string.

unsigned long *f_flag*

A bitmask of flags; the function can set these flags:

- ST_RDONLY — read-only filesystem.
- ST_NOSUID — the filesystem doesn't support **setuid/setgid** semantics.

unsigned long *f_namemax*

The maximum filename length.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EBADF The *fdes* argument isn't an open file descriptor.
- EFAULT The *buf* argument points to an illegal address.
- EINTR A signal was caught during execution.
- EIO An I/O error occurred while reading the filesystem.
- EOVERFLOW One of the values to be returned can't be represented correctly in the structure pointed to by *buf*.

Classification:

fstatvfs() is standard Unix; *fstatvfs64()* is for large-file support

Safety

- Cancellation point No
- Interrupt handler No
- Signal handler Yes
- Thread Yes

Caveats:

The values returned for *f_files*, *f_ffree*, and *f_favail* might not be valid for NFS-mounted filesystems.

See also:

chmod(), *chown()*, *creat()*, *dup()*, *fcntl()*, *link()*, *mknod()*, *open()*,
pipe(), *read()*, *statvfs()*, *statvfs64()*, *time()*, *unlink()*, *utime()*, *write()*

fsync()

© 2004, QNX Software Systems Ltd.

Synchronize the file state

Synopsis:

```
#include <unistd.h>

int fsync( int filedes );
```

Arguments:

filedes The descriptor for the file that you want to synchronize.

Library:

`libc`

Description:

The *fsync()* function forces all queued I/O operations for the file specified by the *filedes* file descriptor to finish, synchronizing the file's state.

Although similar to *fdatasync()*, *fsync()* also guarantees the integrity of file information such as access and modification times.

Returns:

0 for success, or -1 if an error occurs (*errno* is set).

Errors:

EBADF	The <i>filedes</i> argument isn't a valid file descriptor open for writing.
EINVAL	The implementation doesn't support synchronized I/O for the given file.
ENOSYS	The <i>fsync()</i> function isn't supported for the filesystem specified by <i>filedes</i> .

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

aio_fsync(), close(), fcntl(), fdatsync(), open(), read(), sync(), write()

ftell(), ftello()

© 2004, QNX Software Systems Ltd.

Return the current position of a stream

Synopsis:

```
#include <stdio.h>

long int ftell( FILE* fp );

off_t ftello( FILE* fp );
```

Arguments:

fp The stream that you want to get the current position of.

Library:

`libc`

Description:

The *ftell()* function returns the current position of the stream specified by *fp*. This position defines the character that will be read or written by the next I/O operation on the file. You can use the value returned by *ftell()* in a subsequent call to *fseek()* to restore the file position to a previous value.

The *ftello()* function is similar to *ftell()*, except that the position is returned as an `off_t`.

Returns:

The current position of the file or `-1L` if an error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <stdlib.h>

long filesize( FILE *fp )
{
    long int save_pos;
    long size_of_file;
```

```
    /* Save the current position. */
    save_pos = ftell( fp );

    /* Jump to the end of the file. */
    fseek( fp, 0L, SEEK_END );

    /* Get the end position. */
    size_of_file = ftell( fp );

    /* Jump back to the original position. */
    fseek( fp, save_pos, SEEK_SET );

    return( size_of_file );
}

int main( void )
{
    FILE *fp;

    fp = fopen( "file", "r" );

    if( fp != NULL ) {
        printf( "File size=%ld\n", filesize( fp ) );
        fclose( fp );

        return EXIT_SUCCESS;
    }

    return EXIT_FAILURE;
}
```

Classification:

ftell() is ANSI; *ftello()* is standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, fgetpos(), fopen(), fsetpos(), fseek()

Synopsis:

```
#include <sys/timeb.h>

int ftime( struct timeb * timeptr );
```

Arguments:

timeptr A pointer to a **timeb** structure where the function can store the current time; see below.

Library:

libc

Description:

The *ftime()* function stores the current time in the *timeptr* structure. The **timeb** structure contains the following fields:

time_t *time* Time, in seconds, since the Unix Epoch, 00:00:00 January 1, 1970 Coordinated Universal Time (UTC).

unsigned short *millitm* Milliseconds.

short *timezone* Difference in minutes of the timezone from UTC.

short *dstflag* Nonzero if in daylight savings time.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Examples:

```
#include <stdio.h>
#include <time.h>
#include <sys/timeb.h>
#include <stdlib.h>

int main( void )
{
    struct timeb timebuf;
    char *now;

    ftime( &timebuf );
    now = ctime( &timebuf.time );

    /* Note that we're cutting "now" off
     * after 19 characters to avoid the
     * \n that ctime() appends to the
     * formatted time string.
     */

    printf( "The time is %.19s.%hu\n",
           now, timebuf.millitm );

    return EXIT_SUCCESS;
}
```

Produces output similar to the following:

```
The time is Mon Jul 05 15:58:42.870
```

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

*asctime(), clock(), ctime(), difftime(), gmtime(), localtime(), mktime(),
strftime(), time(), tzset()*

ftruncate()*, *ftruncate64()

© 2004, QNX Software Systems Ltd.

Truncate a file

Synopsis:

```
#include <unistd.h>

int ftruncate( int fdes,
               off_t length );

int ftruncate64( int fdes,
                 off64_t length );
```

Arguments:

fdes The descriptor for the file that you want to truncate.

length The length that you want the file to be, in bytes.

Library:

libc

Description:

These functions cause the file referenced by *fdes* to have a size of *length* bytes. If the size of the file previously exceeded *length*, the extra data is discarded (this is similar to using the `F_FREESP` option with *fcntl()*). If the size of the file was previously shorter than *length*, the file size is extended with NUL characters (similar to the `F_ALLOCS` option to *fcntl()*).

The value of the seek pointer isn't modified by a call to *ftruncate()*.

Upon successful completion, the *ftruncate()* function marks the *st_ctime* and *st_mtime* fields of the file for update. If the *ftruncate()* function is unsuccessful, the file is unaffected.

Returns:

Zero for success, or -1 if an error occurred (*errno* is set).

Errors:

EBADF	The <i>fdes</i> argument isn't a valid file descriptor.
EFBIG	The file is a regular file and <i>length</i> is greater than the offset maximum associated with the file.
EINTR	A signal was caught during the call to <i>ftruncate()</i> .
EINVAL	The <i>fdes</i> argument doesn't refer to a file on which this operation is possible, the <i>fdes</i> argument isn't open for writing or the <i>length</i> argument is less than the minimum file size for the specified filesystem.
EIO	An I/O error occurred while reading from or writing to the filesystem.
ENOSYS	The <i>ftruncate()</i> function isn't implemented for the filesystem specified by <i>fdes</i> .
ENOTSUP	The <i>ftruncate()</i> function is implemented for the specified filesystem, but the specific operation (F_ALLOCSP or F_FREESP; see <i>fctl()</i>) isn't supported.
EROFS	The file resides on a read-only filesystem.

Classification:

ftruncate() is POSIX 1003.1; *ftruncate64()* is for large-file support

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

mmap(), *open()*, *shm_open()*, *truncate()*

Synopsis:

```
#include <stdio.h>

int ftrylockfile( FILE* file );
```

Arguments:

file A pointer to the **FILE** object for the file you want to lock.

Library:

`libc`

Description:

The *ftrylockfile()* function is used by a thread to acquire ownership of a **FILE** if the object is available; *ftrylockfile()* is a nonblocking version of *flockfile()*.

Returns:

0 Success.

Nonzero The lock can't be acquired.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

flockfile(), *getc_unlocked()*, *getchar_unlocked()*, *putc_unlocked()*,
putchar_unlocked()

Synopsis:

```
#include <ftw.h>

int ftw( const char *path,
        int (*fn)( const char *fname,
                  const struct stat *sbuf,
                  int flags) ,
        int depth );
```

Arguments:

- path* The path of the directory whose file tree you want to walk.
- fn* A pointer to a function that you want to call for each file; see below.
- depth* The maximum number of file descriptors that *ftw()* can use. The *ftw()* function uses one file descriptor for each level in the tree.
- If *depth* is zero or negative, the effect is the same as if it were 1. The *depth* must not be greater than the number of file descriptors currently available for use. The *ftw()* function is faster if *depth* is at least as large as the number of levels in the tree.

Library:

libc

Description:

The *ftw()* function recursively descends the directory hierarchy identified by *path*. For each object in the hierarchy, *ftw()* calls the user-defined function *fn()*, passing to it:

- a pointer to a NULL-terminated character string containing the name of the object

- a pointer to a `stat` structure (see `stat()`) containing information about the object
- an integer. Possible values of the integer, defined in the `<ftw.h>` header, are:

<code>FTW_F</code>	The object is a file.
<code>FTW_D</code>	The object is a directory.
<code>FTW_DNR</code>	The object is a directory that can't be read. Descendents of the directory aren't processed.
<code>FTW_NS</code>	The <code>stat()</code> failed on the object because the permissions weren't appropriate, or the object is a symbolic link that points to a nonexistent file. The <code>stat</code> buffer passed to <code>fn()</code> is undefined.

The `ftw()` function visits a directory before visiting any of its descendents.

The tree traversal continues until the tree is exhausted, an invocation of `fn()` returns a nonzero value, or some error is detected within `ftw()` (such as an I/O error). If the tree is exhausted, `ftw()` returns zero. If `fn()` returns a nonzero value, `ftw()` stops its tree traversal and returns whatever value was returned by `fn()`.

When `ftw()` returns, it closes any file descriptors it opened; it doesn't close any file descriptors that may have been opened by `fn()`.

Returns:

- 0 Success.
- 1 An error (other than `EACCESS`) occurred (`errno` is set).

Classification:

Standard Unix, `ftw64()` is for large-file support

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Because *ftw()* is recursive, it might terminate with a memory fault when applied to very deep file structures.

This function uses *malloc()* to allocate dynamic storage during its operation. If *ftw()* is forcibly terminated, for example if *longjmp()* is executed by *fn()* or an interrupt routine, *ftw()* doesn't have a chance to free that storage, so it remains permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn()* return a nonzero value at its next invocation.

See also:

longjmp(), *malloc()*, *nftw()*, *stat()*

funlockfile()

© 2004, QNX Software Systems Ltd.

Release ownership of a file

Synopsis:

```
#include <stdio.h>

void funlockfile( FILE* file );
```

Arguments:

file A pointer to the **FILE** object for the file you want to unlock.

Library:

`libc`

Description:

The *funlockfile()* function is used to release ownership of *file* granted to the thread. The behavior is undefined if a thread other than the current owner calls the *funlockfile()* function.

For more information, see *flockfile()*.

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

flockfile(), *ftrylockfile()*, *getc_unlocked()*, *getchar_unlocked()*,
putc_unlocked(), *putchar_unlocked()*

futime()

© 2004, QNX Software Systems Ltd.

Record the modification time for a file

Synopsis:

```
#include <utime.h>

int futime( int fdes,
            const struct utimbuf *times );

struct utimbuf {
    time_t  actime;    /* access time */
    time_t  modtime;  /* modification time */
};
```

Arguments:

- fdes* The descriptor for the file whose modification time you want to get or set.
- times* NULL, or a pointer to a `utimbuf` structure where the function can store the modification time.

Library:

`libc`

Description:

The `futime()` function records the modification time for the file or directory with the descriptor, *fdes*.

If the *times* argument is NULL, the access and modification times of the file or directory are set to the current time. The effective user ID of the process must match the owner of the file or directory, or the process must have write permission to the file or directory, or appropriate privileges in order to use the `futime()` function in this way.

If the *times* argument isn't NULL, it's interpreted as a pointer to a `utimbuf` structure, and the access and modification times of the file or directory are set to the values contained in the *actime* and *modtime* fields in this structure. Only the owner of the file or directory, and processes with appropriate privileges are permitted to use the `futime()` function in this way.

Returns:

- 0 Success.
- 1 An error occurred (*errno*) is set.

Errors:

- EACCES Search permission is denied for a component of *path*, or the *times* argument is NULL, and the effective user ID of the process doesn't match the owner of the file, and write access is denied.
- ENAMETOOLONG
The argument *path* exceeds PATH_MAX in length, or a pathname component is longer than NAME_MAX.
- ENOENT The specified *path* doesn't exist, or *path* is an empty string.
- ENOTDIR A component of *path* isn't a directory.
- EPERM The *times* argument isn't NULL, and the calling process's effective user ID has write access to the file but doesn't match the owner of the file, and the calling process doesn't have the appropriate privileges.
- EROFS The named file resides on a read-only filesystem.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

errno, utime()

Synopsis:

```
#include <wchar.h>

int fwide( FILE * fp,
          int mode );
```

Arguments:

- fp* The stream whose orientation you want to set.
- mode* The orientation mode:
- If *mode* is greater than zero and the stream orientation hasn't been set, *fwide()* flags the stream as wide-oriented.
 - If *mode* is less than zero, *fwide()* behaves similarly, but flags the stream as byte-oriented.
 - If *mode* is zero, *fwide()* returns the stream type without altering the stream.

Library:

`libc`

Description:

The *fwide()* function sets or determines the orientation of the stream *fp*.

Returns:

- > 0 The stream is (now) wide-oriented.
- 0 The stream is unbound.
- < 0 The stream is (now) byte-oriented.

Errors:

EBADF The *fp* argument isn't valid.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

Synopsis:

```
#include <wchar.h>

int fwprintf( FILE * fp,
              const wchar_t * format,
              ... );
```

Arguments:

fp The stream to which you want to send the output.

format A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

Library:

`libc`

Description:

The *fwprintf()* function writes output to the stream specified by *fp*, under control of the *format* specifier.

The *fwprintf()* function is the wide-character version of *fprintf()*.

Returns:

The number of wide characters written, excluding the terminating NUL, or a negative number if an error occurred (*errno* is set).

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*, *vfprintf()*,
vfwprintf(), *vprintf()*, *vsprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*,
wprintf()

Synopsis:

```
#include <stdio.h>

size_t fwrite( const void* buf,
               size_t size,
               size_t num,
               FILE* fp );
```

Arguments:

buf A pointer to a buffer that contains the elements that you want to write.

size The size of each element to write.

num The number of elements to write.

fp The stream to which to write the elements.

Library:

libc

Description:

The *fwrite()* function writes *num* elements of *size* bytes each to the stream specified by *fp*.

Returns:

The number of complete elements successfully written; if an error occurs, this is less than *num*.

Errors:

If an error occurs, *errno* is set to indicate the type of error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

struct student_data {
    int student_id;
    unsigned char marks[10];
};

int main( void )
{
    FILE *fp;
    struct student_data std;
    int i;

    fp = fopen( "file", "w" );
    if( fp != NULL ) {
        std.student_id = 1001;

        for( i = 0; i < 10; i++ ) {
            std.marks[i] = (unsigned char)(85 + i);
        }

        /* write student record with marks */
        i = fwrite( &std, sizeof( struct student_data ), 1, fp );
        printf( "Successfully wrote %d records\n", i );

        fclose( fp );

        if( i == 1 ) {
            return EXIT_SUCCESS;
        }
    }

    return EXIT_FAILURE;
}
```

Classification:

ANSI

Safety

Cancellation point Yes

continued...

Safety

Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, ferror(), fopen()

fwscanf()

© 2004, QNX Software Systems Ltd.

Scan wide-character input from a stream

Synopsis:

```
#include <wchar.h>

int fwscanf( FILE * fp,
             const wchar_t * format,
             ... );
```

Arguments:

fp The stream that you want to read from.

format A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

Library:

`libc`

Description:

The *fwscanf()* function scans input from the stream specified by *fp*, under control of the argument *format*. Following the format string is a list of addresses to receive values.

The *fwscanf()* function is the wide-character version of *fscanf()*.

Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning reached the end of the input stream before storing any values.

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *fscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *vfscanf()*, *vfwscanf()*,
vscanf(), *vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

gai_strerror()

© 2004, QNX Software Systems Ltd.

Return the string associated with a `getaddrinfo()` error code

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netdb.h>

const char * gai_strerror( int ecode );
```

Arguments:

ecode The error code number from the `getaddrinfo()` function.

Library:

`libsocket`

Description:

The `gai_strerror()` function returns a string describing the error code from the `getaddrinfo()` function. Nonzero error codes are defined in `<netdb.h>` as follows:

EALADDRFAMILY	The address family for <i>nodename</i> isn't supported.
EALAGAIN	There was a temporary failure in name resolution.
EALBADFLAGS	Invalid value for <i>ai_flags</i> .
EALFAIL	Nonrecoverable failure in name resolution.
EALFAMILY	The <i>ai_family</i> isn't supported.
EALMEMORY	Memory allocation failure.
EALNODATA	No address associated with the <i>nodename</i> .
EALNONAME	Either the <i>nodename</i> or the <i>servname</i> argument wasn't provided or isn't known.

EALSERVICE	The <i>servname</i> argument isn't supported for <i>ai_socktype</i> .
EALSOCKTYPE	The <i>ai_socktype</i> isn't supported.
EALSYSTEM	System error returned in <i>errno</i> .

Returns:

If called with a proper *ecode* argument, a pointer to a string describing the given error code. If the argument isn't one of the EAL_* values, a pointer to a string whose contents indicate an unknown error.



Don't modify the strings that this function returns.

Classification:

POSIX 1003.1-2001

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

`addrinfo`, `freaddrinfo()`, `getaddrinfo()`

gamma()*, *gamma_r()*, *gammaf()*, *gammaf_r() © 2004, QNX

Software Systems Ltd.

Log gamma function

Synopsis:

```
#include <math.h>

double gamma( double x );

double gamma_r( double x,
                int* signgam );

float gammaf( float x );

float gammaf_r( float x,
                int* signgam );
```

Arguments:

x An arbitrary number.

signgam (*gamma_r()*, *gammaf_r()* only) A pointer to a location where the function can store the sign of $\Gamma(x)$.

Library:

libm

Description:

The *gamma()* and *gamma_r()* functions return the natural log (**ln**) of the *gamma()* function and are equivalent to *lgamma()*. These functions return **ln** | $\Gamma(x)$ |, where $\Gamma(x)$ is defined as follows:

$$\int_0^{\infty} e^{-t} t^{x-1} dt$$

For $x > 0$:

$$\text{For } x < 1: \quad n / (\Gamma(1-x) * \sin(\pi x))$$

The results converge when x is between 0 and 1. The Γ function has the property:

$$\Gamma(N) = \Gamma(N-1) \times N$$

The *gamma** functions compute the log because the Γ function grows very quickly.

The *gamma()* and *gammaf()* functions use the external integer *signgam* to return the sign of $\Gamma(x)$, while *gamma_r()* and *gammaf_r()* use the user-allocated space addressed by *signgamp*.



The *signgam* variable isn't set until *gamma()* or *gammaf()* returns. For example, don't use the expression:

```
g = signgam * exp( gamma( x ) );
```

to compute $g = \Gamma(x)'$. Instead, compute *gamma()* first:

```
lg = gamma(x);
g = signgam * exp( lg );
```

Note that $\Gamma(x)$ must overflow when x is large enough, underflow when $-x$ is large enough, and generate a division by 0 exception at the singularities x a nonpositive integer.

Returns:

$\ln |\Gamma(x)|$



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

lgamma()

Synopsis:

```
#include <sys/socket.h>
#include <netdb.h>

int getaddrinfo( const char * nodename,
                 const char * servname,
                 const struct addrinfo * hints,
                 struct addrinfo ** res );
```

Arguments:

<i>nodename</i>	The node name. A non-NULL <i>nodename</i> may be either a node name or a numeric host address string (i.e. a dotted-decimal IPv4 address or an IPv6 hex address.)
<i>servname</i>	The server name. A non-NULL <i>servname</i> may be either a server name or a decimal port number.
<i>hints</i>	A pointer to an addrinfo structure that provides hints about the type of socket you're supporting. See "Using the <i>hints</i> argument" for more information.
<i>res</i>	The address of a location where the function can store a pointer to a linked list of one or more addrinfo structures.

Library:

libsocket

Description:

The *getaddrinfo()* function performs the functionality of *gethostbyname()* and *getservbyname()* but in a more sophisticated manner.

The *nodename* and *servname* arguments are either pointers to null-terminated strings or NULL. One or both of these two arguments

must be a non-NULL pointer. Normally, a client scenario specifies both *nodename* and *servname*.

On success, the *getaddrinfo()* function stores, in the location pointed to by *res*, a pointer to a linked list of one or more **addrinfo** structures. You can process each **addrinfo** structure in this list by following the *ai_next* pointer until reaching a NULL pointer. Each **addrinfo** structure contains the corresponding *ai_family*, *ai_socktype*, and *ai_protocol* arguments for a call to the *socket()* function. The *ai_addr* argument of the **addrinfo** structure points to a filled-in socket address structure with a length specified by the *ai_addrlen* argument.

Using the *hints* argument

You can optionally pass an **addrinfo** structure, pointed to by the *hints* argument, that provides hints concerning the type of socket that your application supports.

In this structure, all members — except *ai_flags*, *ai_family*, *ai_socktype*, and *ai_protocol* — must be zero or a NULL pointer. The **addrinfo** structure of the *hints* argument can accept various types of sockets:

To accept:	Set:	To:
Any protocol family	<i>ai_family</i>	PF_UNSPEC
Any socket type	<i>ai_socktype</i>	0
Any protocol	<i>ai_protocol</i>	0
All of the above (as well as setting <i>ai_flags</i> to 0)	<i>hints</i>	NULL

The *hints* argument defaults to all possibilities, but you can also use it to limit choices:

- If the application handles only TCP but not UDP, you could set the *ai_socktype* member of the *hints* structure to `SOCK_STREAM`.
- If the application handles only IPv4 but not IPv6, you could set the *ai_family* member of the *hints* structure to `PF_INET`.

Using the *ai_flags* argument in the *hints* structure

You can set the *ai_flags* argument to further configure the *hints* structure. Settings for *ai_flags* include:

AI_PASSIVE Set this bit if you plan to use the returned **addrinfo** structure in a call to *bind()*. In this call, if the *nodename* argument is a NULL pointer, then the IP address portion of the socket address structure *ai_addr* is set to `INADDR_ANY` for an IPv4 address or `IN6ADDR_ANY_INIT` for an IPv6 address.

If you don't set the `AI_PASSIVE` flag, you can use the returned **addrinfo** structure in a call to:

- *connect()* — connectionless or connection-oriented protocol
- *sendto()* — connectionless protocol
- *sendmsg()* — connectionless protocol

In this case, if the *nodename* argument is a NULL pointer, then the IP address portion of the socket address structure *ai_addr* is set to the loopback address.

AI_CANONNAME

Set this bit if you want the *ai_canonname* argument of the first **addrinfo** structure to point to a null-terminated string containing the canonical name of the specified *nodename*.

AI_NUMERICHOST

Set this bit if you want to prevent any type of name resolution service (such as DNS) from being used.

A non-NULL *nodename* string must be a numeric host address string; otherwise, *getaddrinfo()* returns `EAI_NONAME`.

Pitfalls

The arguments to *getaddrinfo()* must be sufficiently consistent and unambiguous or this function will return an error. Here are some problems you may encounter:

- Inconsistent *hints* — for Internet address families, specifying `SOCK_STREAM` for *ai_socktype* while specifying `IPPROTO_UDP` for *ai_protocol*.
- Inconsistent *servname* — specifying a *servname* that's defined only for certain *ai_socktype* values, such as the TFTP service (a datagram service `SOCK_DGRAM`) on `SOCK_STREAM`.
- Undefined service names — specifying a *servname* while specifying `SOCK_RAW` for *ai_socktype*. (Service names aren't defined for the internet `SOCK_RAW` space.)
- Incomplete specifications — specifying a numeric *servname* while leaving *ai_socktype* and *ai_protocol* unspecified. The *getaddrinfo()* function isn't allowed to *glob()* the argument when a numeric *servname* doesn't have a specified socket type.



The *getaddrinfo()* function dynamically allocates space for the following:

- *addrinfo* structures
- socket address structures
- canonical node name strings pointed to by the *addrinfo* structures.

Use *freeaddrinfo()* to free the *addrinfo* structures, and *gai_strerror()* to decipher error codes.

Returns:

Zero for success, or nonzero if an error occurs.

Errors:

To get an explanation of any error code, use *gai_strerror()*.

Examples:

The following code tries to connect to **www.kame.net** service HTTP using a stream socket. It loops through all the addresses available, regardless of the address family. If the destination resolves to an IPv4 address, it uses a AF_INET socket. Similarly, it uses an AF_INET6 socket if it resolves to IPv6. Note that there aren't any hardcoded references to any particular address family; the code works even if *getaddrinfo()* returns addresses that aren't IPv4/v6.

```
struct addrinfo hints, *res, *res0;
int error;
int s;
const char *cause = NULL;

memset(&hints, 0, sizeof(hints));
hints.ai_family = PF_UNSPEC;
hints.ai_socktype = SOCK_STREAM;
error = getaddrinfo("www.kame.net", "http", &hints, &res0);
if (error) {
    errl(1, "%s", gai_strerror(error));
    /*NOTREACHED*/
}
```

```

}
s = -1;
for (res = res0; res; res = res->ai_next) {
    s = socket(res->ai_family, res->ai_socktype,
              res->ai_protocol);
    if (s < 0) {
        cause = "socket";
        continue;
    }

    if (connect(s, res->ai_addr, res->ai_addrlen) < 0) {
        cause = "connect";
        close(s);
        s = -1;
        continue;
    }

    break; /* okay we got one */
}
if (s < 0) {
    err(1, cause);
    /*NOTREACHED*/
}
freeaddrinfo(res0);

```

The following example tries to open a wildcard-listening socket onto the HTTP service for all of the available address families:

```

struct addrinfo hints, *res, *res0;
int error;
int s[MAXSOCK];
int nsock;
const char *cause = NULL;

memset(&hints, 0, sizeof(hints));
hints.ai_family = PF_UNSPEC;
hints.ai_socktype = SOCK_STREAM;
hints.ai_flags = AI_PASSIVE;
error = getaddrinfo(NULL, "http", &hints, &res0);
if (error) {
    err(1, "%s", gai_strerror(error));
    /*NOTREACHED*/
}
nsock = 0;
for (res = res0; res && nsock < MAXSOCK; res = res->ai_next) {
    s[nsock] = socket(res->ai_family, res->ai_socktype,
                    res->ai_protocol);
    if (s[nsock] < 0) {
        cause = "socket";
    }
}

```

```
        continue;
    }

    if (connect(s[nsock], res->ai_addr, res->ai_addrlen) < 0) {
        cause = "connect";
        close(s[nsock]);
        continue;
    }

    nsock++;
}
if (nsock == 0) {
    err(1, cause);
    /*NOTREACHED*/
}
freeaddrinfo(res0);
```

Classification:

POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

addrinfo, *freeaddrinfo()*, *gai_strerror()*

getc()

© 2004, QNX Software Systems Ltd.

Get the next character from a file

Synopsis:

```
#include <stdio.h>

int getc( FILE* fp );
```

Arguments:

fp The stream you want to get the character from.

Library:

libc

Description:

The *getc()* macro gets the next character from the stream designated by *fp*. The character is returned as an `int` value.

Returns:

The next character from the stream *fp*, cast as `(int) (unsigned char)`, or EOF if an end-of-file or error condition occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE* fp;
    int c;

    fp = fopen( "file", "r" );
    if( fp != NULL ) {
        while( ( c = getc( fp ) ) != EOF ) {
            putchar(c);
        }
    }
}
```

```
    }  
    fclose( fp );  
    return EXIT_SUCCESS;  
}  
return EXIT_FAILURE;  
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

getc() is a macro.

See also:

errno, *feof()*, *ferror()*, *fgetc()*, *fgetchar()*, *fgets()*, *fopen()*, *getchar()*,
gets(), *putc()*, *putc_unlocked()*, *putchar()*, *putchar_unlocked()*,
ungetc()

getc_unlocked()

© 2004, QNX Software Systems Ltd.

Get the next character from a file

Synopsis:

```
#include <stdio.h>

int getc_unlocked( FILE *fp );
```

Arguments:

fp The stream you want to get the character from.

Library:

`libc`

Description:

The *getc_unlocked()* function is a thread-unsafe version of *getc()*. You can use it safely only when the invoking thread has locked *fp* using *flockfile()* (or *ftrylockfile()*) and *funlockfile()*.

Returns:

The next character from the input stream pointed to by *fp*, or EOF if an end-of-file or error condition occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Classification:

POSIX 1003.1

Safety

Cancellation point Yes

Interrupt handler No

continued...

Safety

Signal handler	No
Thread	No

See also:

feof(), *ferror()*, *flockfile()*, *getc()*, *getchar()*, *getchar_unlocked()*,
putc(), *putc_unlocked()*, *putchar()*, *putchar_unlocked()*

getchar()

© 2004, QNX Software Systems Ltd.

Get a character from stdin

Synopsis:

```
#include <stdio.h>

int getchar( void );
```

Library:

libc

Description:

The *getchar()* function is equivalent to *getc()* on the *stdin* stream.

Returns:

The next character from the input stream pointed to by *stdin*, cast as (int) (unsigned char), or EOF if an end-of-file or error condition occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    FILE *fp;
    int c;

    /* Get characters from "file" instead of
     * stdin.
     */
    fp = freopen( "file", "r", stdin );
    while( ( c = getchar() ) != EOF ) {
        putchar(c);
    }

    fclose( fp );
}
```

```
        return EXIT_SUCCESS;
    }
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *feof()*, *ferror()*, *fgetc()*, *fgetchar()*, *getc()*, *putc()*,
putc_unlocked(), *putchar()*, *putchar_unlocked()*

getchar_unlocked()

© 2004, QNX Software Systems Ltd.

Get a character from stdin

Synopsis:

```
#include <stdio.h>

int getchar_unlocked( void );
```

Library:

libc

Description:

The *getchar_unlocked()* function is a thread-unsafe version of *getchar()*. You can use it safely only when the invoking thread has locked *stdin* using *flockfile()* (or *ftrylockfile()*) and *funlockfile()*.

Returns:

The next character from the input stream pointed to by *stdin*, cast as (**int**) (**unsigned char**), or EOF if an end-of-file or error condition occurs (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

*feof(), ferror(),getc(),getc_unlocked(),getchar(),putc(),
putc_unlocked(), putchar(), putchar_unlocked()*

getcwd()

© 2004, QNX Software Systems Ltd.

Get the name of the current working directory

Synopsis:

```
#include <unistd.h>

char* getcwd( char* buffer,
              size_t size );
```

Arguments:

buffer A pointer to a buffer where the function can store the directory name.

size The size of the buffer, in bytes.

Library:

`libc`

Description:

The *getcwd()* function returns the name of the current working directory. *buffer* is a pointer to a buffer of at least *size* bytes where the NUL-terminated name of the current working directory will be placed.

The maximum size that might be required for *buffer* is `PATH_MAX + 1` bytes. See `<limits.h>`.

Returns:

The address of the string containing the name of the current working directory, or `NULL` if an error occurs (*errno* is set).

Errors:

`EINVAL` The argument *size* is negative or 0.

`ELOOP` Too many levels of symbolic links.

`ENOSYS` The *getcwd()* function isn't implemented for the filesystem specified in the current working directory.

ERANGE The buffer is too small (as specified by *size*) to contain the name of the current working directory.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <limits.h>

int main( void )
{
    char* cwd;
    char buff[PATH_MAX + 1];

    cwd = getcwd( buff, PATH_MAX + 1 );
    if( cwd != NULL ) {
        printf( "My working directory is %s.\n", cwd );
    }

    return EXIT_SUCCESS;
}
```

produces the output:

```
My working directory is /home/bill.
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

Caveats:

There is only one current working directory per *process*. In a multithreaded application, any thread calling *chdir()* will change the current working directory for *all threads* in that process.

See also:

chdir(), *errno*, *mkdir()*, *rmdir()*

Synopsis:

```
#include <unistd.h>

int getdomainname( char * name,
                  size_t namelen );
```

Arguments:

name A buffer where the function can store the domain name.

namelen The size of the name array.

Library:

libsocket

Description:

The *getdomainname()* function gets the standard domain name for the current processor and stores it in the buffer that *name* points to. The name is null-terminated.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

EFAULT The *name* or *namelen* parameters gave an invalid address.

Classification:

Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

setdomainname()

Synopsis:

```
#include <unistd.h>

int  getdtablesize( void );
```

Library:

libc

Description:

Each process has a fixed size descriptor table, which is guaranteed to have at least 20 slots. The entries in the descriptor table are numbered with small integers starting at 0. The *getdtablesize()* returns the size of this table.

This function is equivalent to *getrlimit()* with the RLIMIT_NOFILE option.

Returns:

The size of the file descriptor table.

Classification:

Legacy Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

close(), dup(), getrlimit(), open(), select(), sysconf()

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

gid_t getegid( void );
```

Library:

libc

Description:

The *getegid()* function gets the effective group ID for the calling process.

Returns:

The calling process's effective group ID. This function can't fail.

Examples:

```
/*
 * Print the effective group ID of a process
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main( void )
{
    printf( "My effective group ID is %d\n", getegid() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

geteuid(), getgid(), getuid(), setegid()

Synopsis:

```
#include <stdlib.h>

char* getenv( const char* name );
```

Arguments:

name The name of the environment variable whose value you want to get.

Library:

libc

Description:

The *getenv()* function searches the environment list for a string in the form *name=value* and returns a pointer to a string containing the *value* for the specified *name*. The matching is case-sensitive.

Returns:

A pointer to the *value* assigned to *name*, or NULL if *name* wasn't found in the environment.



Don't modify the returned string.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    char* path;

    path = getenv( "INCLUDE" );
    if( path != NULL ) {
        printf( "INCLUDE=%s\n", path );
        return EXIT_SUCCESS;
    }
}
```

```
        return EXIT_FAILURE;  
    }
```

Classification:

ANSI, POSIX 1003.1a

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

The *getenv()* function manipulates the environment pointed to by the global *environ* variable.

See also:

clearenv(), *environ*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *putenv()*, *searchenv()*, *setenv()*, *spawn*()* functions, *system()*

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

uid_t geteuid( void );
```

Library:

libc

Description:

The *geteuid()* function gets the effective user ID for the calling process.

Returns:

The calling process's effective user ID.

Examples:

```
/*
 * Print the effective user ID of a process.
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main( void )
{
    printf( "My effective user ID is %d\n", geteuid() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

getegid(), getgid(), getuid(), seteuid()

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

gid_t getgid( void );
```

Library:

libc

Description:

The *getgid()* function gets the group ID for the calling process.

Returns:

The calling process's group ID. This function can't fail.

Examples:

```
/*
 * Print the group id of a process.
 */

#include <stdlib.h>
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>

int main( void )
{
    printf( "I belong to group ID %d\n", getgid() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

getegid(), geteuid(), getuid()

Synopsis:

```
#include <grp.h>

struct group* getgrent( void );
```

Library:

libc

Description:

The *getgrent()* function returns the next entry from the group database, although no particular order is guaranteed. This function uses a static buffer that's overwritten by each call.



The *getgrent()*, *getgrgid()*, and *getgrnam()* function share the same static buffer.

Returns:

The next entry from the group database. When you first call *getgrent()*, the group database is opened. It remains open until either *getgrent()* returns NULL to signify end-of-file, or you call *endgrent()*.

Errors:

The *getgrent()* function uses the following functions, and as a result, *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

Examples:

```
/*
 * This program loops, reading a group name from
 * standard input and checking to see if it is a valid
 * group. If it isn't valid, the entire contents of the
 * group database are printed.
 */
#include <stdio.h>
#include <stdlib.h>
#include <grp.h>
#include <limits.h>

int main( void )
{
    struct group* gr;
    char    buf[80];

    setgrent();
    while( gets(buf) != NULL) {
        if( (gr=getgrnam(buf)) != (struct group *)0) {
            printf("Valid group is: %s\n",gr->gr_name);
        } else {
            setgrent();
            while( (gr=getgrent()) != (struct group *)0 )
                printf("%s\n",gr->gr_name);
        }
    }
    endgrent();
    return( EXIT_SUCCESS );
}
```

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

endgrent(), errno, getgrgid(), getgrnam(), getpwent(), setgrent()

getgrgid()

© 2004, QNX Software Systems Ltd.

Get information about the group with a given ID

Synopsis:

```
#include <sys/types.h>
#include <grp.h>

struct group* getgrgid( gid_t gid );
```

Arguments:

gid The ID of the group you want to get information about.

Library:

`libc`

Description:

The *getgrgid()* function lets a process gain more knowledge about group *gid*. This function uses a static buffer that's overwritten by each call.



The *getgrent()*, *getgrgid()*, and *getgrnam()* functions share the same static buffer.

Returns:

A pointer to an object of type *struct group* containing an entry from the group database with a matching *gid*. On error or failure to find an entry with a matching *gid*, a NULL pointer is returned.

Examples:

```
/*
 * Print a list of all users in your group
 */
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <grp.h>
```



```
int main( void )
{
    struct group* g;
    char** p;

    if( ( g = getgrgid( getgid() ) ) == NULL ) {
        fprintf( stderr, "getgrgid: NULL pointer\n" );
        return( EXIT_FAILURE );
    }
    printf( "group name:%s\n", g->gr_name );
    for( p = g->gr_mem; *p != NULL; p++ ) {
        printf( "\t%s\n", *p );
    }
    return( EXIT_SUCCESS );
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getgrent(), *getgrgid_r()*, *getgrnam()*

getgrgid_r()

© 2004, QNX Software Systems Ltd.

Get information about the group with a given ID

Synopsis:

```
#include <sys/types.h>
#include <grp.h>

int getgrgid_r ( gid_t gid,
                 struct group* grp,
                 char* buffer,
                 size_t bufsize,
                 struct group** result );
```

Arguments:

<i>gid</i>	The ID of the group you want to get information about.
<i>grp</i>	A pointer to a group structure where the function can store information about the group.
<i>buffer</i>	A buffer from which to allocate any memory required.
<i>bufsize</i>	The size of the buffer.
<i>result</i>	The address of a pointer that <i>getgrgid_r()</i> sets to the same pointer as <i>grp</i> on success, or to NULL if the function can't find the group.

Library:

libc

Description:

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, *getgrgid_r()* updates the group structure pointed by *grp* and stores a pointer to that structure at the location pointed by *result*. The structure contains an entry from the group database with a matching *gid*.

This function allocates storage referenced by the group structure from the memory provided with the *buffer* parameter, which is *bufsize* characters in size. You can determine the maximum size needed for

this *buffer* by calling *sysconf()* with an argument of `_SC_GETGR_R_SIZE_MAX`.

The *getgrgid_r()* stores a NULL pointer at the location pointed by *result* on error or if the requested entry isn't found.

Returns:

Zero for success, or an error number if an error occurred.

Errors:

ERANGE Insufficient storage was supplied via *buffer* and *bufsize* to contain the resulting *group* structure.

The *getgrgid_r()* function uses the following functions, and as a result, *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getgrgid(), getgrnam(), getgrnam_r(), getlogin(), sysconf()

Synopsis:

```
#include <sys/types.h>
#include <grp.h>

struct group* getgrnam( const char* name );
```

Arguments:

name The name of the group you want to get information about.

Library:

`libc`

Description:

The *getgrnam()* function lets a process gain more knowledge about the group named *name*. This function uses a static buffer that's overwritten by each call.



The *getgrent()*, *getgrgid()*, and *getgrnam()* functions share the same static buffer.

Returns:

A pointer to an object of type `struct group` containing an entry from the group database with a matching name, or NULL on error or failure to find an entry with a matching name.

Examples:

```
/*
 * Print the name of all users in the group given in
 * argv[1]
 */
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>
#include <sys/types.h>
#include <grp.h>
```

```
int main( int argc, char** argv )
{
    struct group* g;
    char** p;

    if( ( g = getgrnam( argv[1] ) ) == NULL ) {
        fprintf( stderr, "getgrnam: %s failed\n",
            argv[1] );
        return( EXIT_FAILURE );
    }
    printf( "group name:%s\n", g->gr_name );
    for( p = g->gr_mem; *p != NULL; p++ ) {
        printf( "\t%s\n", *p );
    }
    return( EXIT_SUCCESS );
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:*getgrent(), getgrgid(), getgrnam_r()*

Synopsis:

```
#include <sys/types.h>
#include <grp.h>

int getgrnam_r( const char* name,
                struct group* grp,
                char* buffer,
                size_t bufsize,
                struct group** result );
```

Arguments:

name The name of the group you want to get information about.

grp A pointer to a **group** structure where the function can store information about the group.

buffer A buffer from which to allocate any memory required.

bufsize The size of the buffer.

result The address of a pointer that *getgrgid_r()* sets to the same pointer as *grp* on success, or to NULL if the function can't find the group.

Library:

`libc`

Description:

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, the *getgrnam_r()* function updates the group structure pointed by *grp* and stores a pointer to that structure at the location pointed by *result*. The structure contains an entry from the group database with a matching *name*.

This function allocates storage referenced by the group structure from the memory provided with the *buffer* parameter, which is *bufsize* characters in size. You can determine the maximum size needed for

this *buffer* by calling *sysconf()* with an argument of `_SC_GETGR_R_SIZE_MAX`.

The *getgrnam_r()* stores a NULL pointer at the location pointed by *result* on error or if the requested entry isn't found.

Returns:

Zero for success, or an error number if an error occurred.

Errors:

ERANGE Insufficient storage was supplied via *buffer* and *bufsize* to contain the *group* structure.

The *getgrnam_r()* function uses the following functions, and as a result, *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getrgid(), getrgid_r(), getgrnam(), getlogin(), sysconf()

getgrouplist()

© 2004, QNX Software Systems Ltd.

Determine the group access list for a user

Synopsis:

```
#include <unistd.h>

int getgrouplist( const char *name,
                  gid_t basegid,
                  gid_t *groups,
                  int *ngroups );
```

Arguments:

<i>name</i>	The name of the user.
<i>basegid</i>	The <i>basegid</i> is automatically included in the list of groups. Typically this value is given as the group number from the password file.
<i>groups</i>	A pointer to an array where the function can store the group IDs.
<i>ngroups</i>	A pointer to the size of the <i>groups</i> array. The function sets the value pointed to by <i>ngroups</i> to be the actual number of groups found.

Library:

libc

Description:

The *getgrouplist()* function reads the group file and determines the group access list for the user specified in *name*.

Returns:

-1 if the size of the group list is too small to hold all the user's groups. The group array is filled with as many groups as fit.

Files:

`/etc/group` Group membership list.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *getgrouplist()* function uses the routines based on *getgrent()*. If the invoking program uses any of these routines, the group structure will be overwritten in the call to *getgrouplist()*.

See also:

initgroups(), *setgroups()*

getgroups()

© 2004, QNX Software Systems Ltd.

Get the supplementary group IDs of the calling process

Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

int getgroups( int gidsetsize,
              gid_t grouplist[] );
```

Arguments:

gidsetsize The size of the *grouplist* array.

grouplist An array that the function can fill in with the process's supplementary group IDs.

Library:

libc

Description:

The *getgroups()* function fills the array *grouplist* with the supplementary group IDs of the calling process. The values of array entries with indices greater than or equal to the returned value are undefined.

Returns:

The number of supplementary groups IDs; this value is zero if NGROUPS_MAX is zero. A value of -1 indicates an error (*errno* is set).

Errors:

EINVAL The *gidsetsize* argument isn't equal to zero, and is less than the number of supplementary group IDs.

Examples:

```
/*
 * Print the supplementary group IDs of
 * the calling process.
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main( void )
{
    int    gidsize;
    gid_t  *grouplist;
    int    i;

    gidsize = getgroups( 0, NULL );
    grouplist = malloc( gidsize * sizeof( gid_t ) );
    getgroups( gidsize, grouplist );
    for( i = 0; i < gidsize; i++ )
        printf( "%d\n", ( int ) grouplist[i] );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:*errno, getegid(), geteuid(), getgid(), getuid(), setgroups()*

gethostbyaddr()

© 2004, QNX Software Systems Ltd.

Get a network host entry, given an Internet address

Synopsis:

```
#include <netdb.h>

struct hostent * gethostbyaddr( const void * addr,
                                socklen_t len,
                                int type );
```

Arguments:

addr A pointer to the binary-format (i.e. not NULL-terminated) address in network byte order.

len The length, in bytes, of *addr*.

type The type of address. Currently, this must be AF_INET.

Library:

`libsocket`

Description:

The *gethostbyaddr()* function searches for information associated with a host, which has the address pointed to by *addr* within the address family specified by *type*, opening a connection to the database if necessary.

This function returns a pointer to a structure of type `hostent` that describes an Internet host. This structure contains either the information obtained from a name server, or broken-out fields from a line in `/etc/hosts`.

You can use *sethostent()* to request the use of a connected TCP socket for queries. If the *stayopen* flag is nonzero, all queries to the name server will use TCP and the connection will be retained after each call to *gethostbyaddr()* or *gethostbyname()*. If the *stayopen* flag is zero, queries use UDP datagrams.

Returns:

A pointer to a valid `hostent` structure, or NULL if an error occurs (`h_errno` is set).

Errors:

See *herror()*.

Examples:

Use the *gethostbyaddr()* function to find a host:

```
struct sockaddr_in client;
struct hostent* host;

int sock, fd, len;

:

len = sizeof( client );

fd = accept( sock, (struct sockaddr*)&client, &len );

if( fd == -1 ) {
    perror( "accept" );
    exit( 1 );
}

host = gethostbyaddr( (const void*)&client.sin_addr,
                    sizeof(struct in_addr),
                    AF_INET );

printf( "Connection from %s: (%s)\n",
        host ? host->h_name : "<unknown>",
        inet_ntoa( client.sin_addr ) );

:
```

Files:

`/etc/hosts` Host database file.

`/etc/resolv.conf`

Resolver configuration file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data storage; if you need the data for future use, copy it before any subsequent calls overwrite it. Currently, only the Internet address format is understood.

See also:

endhostent(), *gethostbyname()*, *gethostbyaddr_r()*, *gethostent()*, *herror()*, *hostent*, *sethostent()*

/etc/hosts, */etc/resolv.conf* in the *Utilities Reference*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <netdb.h>

struct hostent * gethostbyaddr_r(
    const void * addr,
    socklen_t length,
    int type,
    struct hostent * result,
    char * buffer,
    int buflen,
    int * h_errnop );
```

Arguments:

<i>addr</i>	A pointer to the binary-format (i.e. not NULL-terminated) address in network byte order.
<i>length</i>	The length, in bytes, of <i>addr</i> .
<i>type</i>	The type of address. Currently, this must be AF_INET.
<i>result</i>	A pointer to a struct hostent where the function can store the host entry.
<i>buffer</i>	A pointer to a buffer that the function can use during the operation to store host database entries; <i>buffer</i> should be large enough to hold all of the data associated with the host entry. A 2K buffer is usually more than enough; a 256-byte buffer is safe in most cases.
<i>buflen</i>	The length of the area pointed to by <i>buffer</i> .
<i>h_errnop</i>	A pointer to a location where the function can store an <i>herrno</i> value if an error occurs.

Library:

`libsocket`

Description:

The *gethostbyaddr_r()* function is a thread-safe version of *gethostbyaddr()*. This function gets the network host entry for the host specified by *addr*. The *addr* argument is the network address of the specified network family, *type*. The buffer for *addr* is at least *length* bytes.

If you need to convert a text-based address into the format necessary for use as *gethostbyaddr_r()*'s *addr*, see *inet_pton()*.

Returns:

A pointer to *result*, or NULL if an error occurs.

Errors:

If an error occurs, the `int` pointed to by *h_errnop* is set to:

- ERANGE The supplied *buffer* isn't large enough to store the result.
- HOST_NOT_FOUND Authoritative answer: Unknown host.
- NO_ADDRESS No address associated with name; look for an MX record.
- NO_DATA Valid name, but no data record of the requested type. The name is known to the name server, but has no IP address associated with it — this isn't a temporary error. Another type of request to the name server using this domain name will result in an answer (e.g. a mail-forwarder may be registered for this domain).

NO_RECOVERY

Unknown server error. An unexpected server failure was encountered. This is a nonrecoverable network error.

TRY_AGAIN

Nonauthoritative answer: Host name lookup failure. This is usually a temporary error and means that the local server didn't receive a response from an authoritative server. A retry at some later time may succeed.

Files:

`/etc/hosts` Local host database file.

`/etc/resolv.conf`
Resolver configuration file.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

gethostbyaddr(), *gethostbyname()*, *gethostbyname_r()*, *inet_ntop()*, *inet_pton()*

`/etc/hosts`, `/etc/resolv.conf` in the *Utilities Reference*

gethostbyname()*, *gethostbyname2()

© 2004, QNX Software

Systems Ltd.

Get a network host entry, given a name

Synopsis:

```
#include <netdb.h>

struct hostent * gethostbyname( const char * name );

struct hostent * gethostbyname2( const char * name,
                                 int af );
```

Arguments:

name The name of the Internet host whose entry you want to find.

af (*gethostbyname2()* only) The address family; one of:

- AF_INET
- AF_INET6

Library:

`libsocket`

Description:

The *gethostbyname()* routine gets the network host entry for a given name. It returns a pointer to a structure of type `hostent` that describes an Internet host. This structure contains either the information obtained from a name server, or broken-out fields from a line in `/etc/hosts`.

When using the name server, *gethostbyname()* searches for the named host in the current domain and in the domain's parents, unless the name ends in a dot.



If the name doesn't contain a dot, and the environment variable **HOSTALIASES** contains the name of an alias file, the alias file is first searched for an alias matching the input name. This file has the same form as */etc/hosts*.

You can use *sethostent()* to request the use of a connected TCP socket for queries. If the *stayopen* flag is nonzero, all queries to the name server use TCP and the connection is retained after each call to *gethostbyname()* or *gethostbyaddr()*. If the *stayopen* flag is zero, queries use UDP datagrams.

The *gethostbyname2()* function is an evolution of the *gethostbyname()* function that lets you look up host names in address families other than AF_INET. If you specify an invalid address family, the function returns NULL and sets *h_errno* to NETDB_INTERNAL.

Returns:

A pointer to a valid **hostent** structure, or NULL if an error occurs (*h_errno* is set).

Errors:

See *herror()*.

Files:

/etc/hosts Host database file.

/etc/resolv.conf
 Resolver configuration file.

For information about these files, see the *Utilities Reference*.

Environment variables:

HOSTALIASES

Name of the alias file that *gethostbyname()* is to search first when the hostname doesn't contain a dot.

Classification:

POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data storage; if you need the data for future use, copy it before any subsequent calls overwrite it.

See also:

endhostent(), *gethostbyaddr()*, *gethostbyname_r()*, *gethostent()*, *herror()*, **hostent**, *sethostent()*

/etc/hosts, */etc/resolv.conf* in the *Utilities Reference*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <netdb.h>

struct hostent * gethostbyname_r( const char * name,
                                  struct hostent * result,
                                  char * buffer,
                                  int buflen,
                                  int * h_errnop );
```

Arguments:

<i>name</i>	The name of the Internet host whose entry you want to find.
<i>result</i>	A pointer to a struct hostent where the function can store the host entry.
<i>buffer</i>	A pointer to a buffer that the function can use during the operation to store host database entries; <i>buffer</i> should be large enough to hold all of the data associated with the host entry. A 2K buffer is usually more than enough; a 256-byte buffer is safe in most cases.
<i>buflen</i>	The length of the area pointed to by <i>buffer</i> .
<i>h_errnop</i>	A pointer to a location where the function can store an <i>errno</i> value if an error occurs.

Library:

libsocket

Description:

The *gethostbyname_r()* function is a thread-safe version of *gethostbyname()*. This function gets the network host entry for the host specified by *name*, and stores the entry in the `struct hostent` pointed to by *result*.

Returns:

A pointer to *result*, or NULL if an error occurs.

Errors:

If an error occurs, the `int` pointed to by *h_errnop* is set to:

ERANGE	The supplied <i>buffer</i> isn't large enough to store the result.
HOST_NOT_FOUND	Authoritative answer: Unknown host.
NO_ADDRESS	No address associated with name; look for an MX record.
NO_DATA	Valid name, but no data record of the requested type. The name is known to the name server, but has no IP address associated with it — this isn't a temporary error. Another type of request to the name server using this domain name will result in an answer (e.g. a mail-forwarder may be registered for this domain).
NO_RECOVERY	Unknown server error. An unexpected server failure was encountered. This is a nonrecoverable network error.
TRY_AGAIN	Nonauthoritative answer: Host name lookup failure. This is usually a temporary error and means that the local server didn't receive a response from an

authoritative server. A retry at some later time may succeed.

Files:

`/etc/hosts` Local host database file.

`/etc/resolv.conf`
Resolver configuration file.

For information about these files, see the *Utilities Reference*.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

gethostbyaddr(), *gethostbyaddr_r()*, *gethostbyname()*

`/etc/hosts`, `/etc/resolv.conf` in the *Utilities Reference*

gethostent()

© 2004, QNX Software Systems Ltd.

Read the next line of the host database file

Synopsis:

```
#include <netdb.h>

struct hostent * gethostent( void );
```

Library:

libsocket

Description:

The *gethostent()* routine reads the next line in the host database file.

Returns:

A pointer to a valid **hostent** structure, or NULL if an error occurs.

Files:

`/etc/hosts` Host database file.
`/etc/resolv.conf`
Resolver configuration file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data storage; if you need the data for future use, copy it before any subsequent calls overwrite it.

Currently, this function understands only the Internet address format.

See also:

endhostent(), *gethostbyaddr()*, *gethostbyname()*, *gethostent_r()*,
hostent, *sethostent()*

/etc/hosts, */etc/resolv.conf* in the *Utilities Reference*

gethostent_r()

© 2004, QNX Software Systems Ltd.

Read the next line of the host database file

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
#include <netdb.h>

struct hostent * gethostent_r( FILE ** hostf,
                               struct hostent * result,
                               char * buffer,
                               int buflen,
                               int * h_errnop );
```

Arguments:

<i>hostf</i>	NULL, or the address of the FILE * pointer associated with the host database file.
<i>result</i>	A pointer to a struct hostent where the function can store the host entry.
<i>buffer</i>	A pointer to a buffer that the function can use during the operation to store host database entries; <i>buffer</i> should be large enough to hold all of the data associated with the host entry. A 2K buffer is usually more than enough; a 256-byte buffer is safe in most cases.
<i>buflen</i>	The length of the area pointed to by <i>buffer</i> .
<i>h_errnop</i>	A pointer to a location where the function can store an <i>herrno</i> value if an error occurs.

Library:

libsocket

Description:

The *gethostent_r()* function is a thread-safe version of the *gethostent()* function. This function gets the local host's entry. If the pointer pointed to by *hostf* is NULL, *gethostent_r()* opens `/etc/hosts` and returns its file pointer in *hostf* for later use. It's the calling process's responsibility to close the host file with *fclose()*.



The first time that you call *gethostent_r()*, pass NULL in the pointer pointed to by *hostf*.

Returns:

A pointer to *result*, or NULL if an error occurs.

Errors:

If an error occurs, the `int` pointed to by *h_errnop* is set to:

ERANGE	The supplied <i>buffer</i> isn't large enough to store the result.
HOST_NOT_FOUND	Authoritative answer: Unknown host.
NO_ADDRESS	No address associated with name, look for an MX record.
NO_DATA	Valid name, no data record of the requested type. The name is known to the name server, but has no IP address associated with it — this isn't a temporary error. Another type of request to the name server using this domain name will result in an answer (e.g. a mail-forwarder may be registered for this domain).
NO_RECOVERY	Unknown server error. An unexpected server failure was encountered. This is a nonrecoverable network error.

TRY_AGAIN Nonauthoritative answer: Host name lookup failure. This is usually a temporary error and means that the local server didn't receive a response from an authoritative server. A retry at some later time may succeed.

Files:

`/etc/hosts` Local host database file.

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

endhostent(), *gethostent()*, *sethostent()*
`/etc/hosts` in the *Utilities Reference*

Synopsis:

```
#include <unistd.h>

int gethostname( char * name,
                size_t namelen );
```

Arguments:

- name* A buffer where the function can store the host name.
- namelen* The size of the buffer.

Library:

libc

Description:

The *gethostname()* function stores in *name* the standard hostname for the current processor, as previously set by *sethostname()*. The parameter *namelen* specifies the size of the *name* array. The returned name is NULL-terminated unless insufficient space is provided.



This function gets the value of the **_CS_HOSTNAME** configuration string, not that of the **HOSTNAME** environment variable.

Returns:

- 0 Success.
- 1 An error occurred (*errno* isn't set).

Classification:

Standard Unix, POSIX 1003.1g (draft)

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Hostnames are limited to MAXHOSTNAMELEN characters (defined in `<sys/param.h>`).

See also:

sethostname()

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <ifaddrs.h>

int getifaddrs( struct ifaddrs ** ifap );
```

Arguments:

ifap The address of a location where the function can store a pointer to a linked list of **ifaddrs** structures that contain the data related to the network interfaces on the local machine.

Library:

libsocket

Description:

The *getifaddrs()* function stores a reference to a linked list of the network interfaces on the local machine in the memory referenced by *ifap*.

The data returned by *getifaddrs()* is dynamically allocated; you should free it by calling *freeifaddrs()* when you no longer need it.

Returns:

0 Success.

-1 An error occurred (*errno* is set).

Errors:

The *getifaddrs()* function may fail and set *errno* for any of the errors specified by:

- *ioctl()*
- *malloc()*

- *socket()*
- *sysctl()*

Classification:

Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *freeifaddrs()*, *ifaddrs*, *ioctl()*, *malloc()*, *socket()*, *sysctl()*

Synopsis:

```
#include <unistd.h>

#define GETIOVBASE( iov ) ...
```

Arguments:

iov The `iov_t` structure from which you want to get the base member.

Library:

`libc`

Description:

This macro evaluates to the *iov_base* member of the given `iov_t` structure.

Returns:

The *iov_base* member of the `iov_t` structure, which is of type `void *`.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

GETIOVLEN(), *SETIOV()*

Synopsis:

```
#include <unistd.h>

#define GETIOVLEN( iov ) ...
```

Arguments:

iov The `iov_t` structure from which you want to get the length member.

Library:

`libc`

Description:

This macro evaluates to the *iov_len* member of the given `iov_t` structure.

Returns:

The *iov_len* member of the `iov_t` structure, which is of type `size_t`.

Classification:

QNX Neutrino

Safety

Cancellation point	No
Interrupt handler	Yes
Signal handler	Yes
Thread	Yes

See also:

GETIOVBASE(), SETIOV()

Synopsis:

```
#include <sys/time.h>

int getitimer ( int which,
               struct itimerval *value );
```

Arguments:

- which* The interval time whose value you want to get. Currently, this must be ITIMER_REAL.
- value* A pointer to a `itimerval` structure where the function can store the value of the interval timer.

Library:

libc

Description:

The system provides each process with several interval timers, defined in `<sys/time.h>`. The `getitimer()` function stores the current value of the timer specified by *which* into the structure pointed to by *value*.

A timer value is defined by the `itimerval` structure (see `gettimeofday()` for the definition of `timeval`), which includes the following members:

```
struct timeval  it_interval;    /* timer interval */
struct timeval  it_value;      /* current value */
```

The *it_value* member indicates the time to the next timer expiration. The *it_interval* member specifies a value to be used in reloading *it_value* when the timer expires. Setting *it_value* to 0 disables a timer, regardless of the value of *it_interval*. Setting *it_interval* to 0 disables a timer after its next expiration (assuming *it_value* is nonzero).

Time values smaller than the resolution of the system clock are rounded up to the resolution of the system clock.

The interval timers include:

ITIMER_REAL Decrements in real time. A SIGALRM signal is delivered when this timer expires.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

EINVAL The specified number of seconds is greater than 100,000,000, the number of microseconds is greater than or equal to 1,000,000, or the *which* argument is unrecognized.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

alarm(), *gettimeofday()*, *pthread_attr_setscope()*, *pthread_sigmask()*, *setitimer()*, *sigprocmask()*, *sleep()*, *sysconf()*

Synopsis:

```
#include <unistd.h>

char* getlogin( void ) ;
```

Library:

libc

Description:

The *getlogin()* function returns a pointer to a string containing the login name of the user associated with the calling process.

Returns:

A pointer to a string containing the user's login name, or NULL if the user's login name can't be found.

The return value from *getlogin()* may point to static data and, therefore, may be overwritten by each call.

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getlogin_r(), getpwnam(), getpwuid()

Synopsis:

```
#include <unistd.h>

int getlogin_r( char* name,
               size_t namesize );
```

Arguments:

name A buffer where the function can store the user name.

namesize The size of the buffer.

Library:

`libc`

Description:

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, the *getlogin_r()* function puts the login name of the user associated with the calling process in the character array pointed to by *name*. The array is *namesize* characters long and should have space for the name and the terminating NULL character. The maximum size of the login name is `_POSIX_LOGIN_NAME_MAX`.

If *getlogin_r()* is successful, *name* points to the name the user used at login, even if there are several login names with the same user ID.

Returns:

EOK Success.

ERANGE Insufficient storage was supplied via the *name* and *namesize* arguments to contain the user's name.

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getlogin(), *getpwnam()*, *getpwnam_r()*, *getpwuid()*, *getpwuid_r()*

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>
#include <netdb.h>

int getnameinfo(const struct sockaddr *sa,
                socklen_t salen,
                char *host, size_t hostlen,
                char *serv, size_t servlen,
                int flags);
```

Arguments:

<i>sa</i>	Points to either a sockaddr_in structure (for IPv4) or a sockaddr_in6 structure (for IPv6) that holds the IP address and port number.
<i>salen</i>	Length of the sockaddr_in or sockaddr_in6 structure.
<i>host</i>	Buffer pointer for the host.
<i>hostlen</i>	Size of the host buffer.
<i>serv</i>	Buffer pointer for the server.
<i>servlen</i>	Length of the server buffer.
<i>flags</i>	Change the default action of <i>getnameinfo()</i> . By default, the fully qualified domain name (FQDN) for the host is looked up in the DNS and returned.

These flags are defined in **<netdb.h>**:

NL_NOFQDN	Only the nodename portion of the FQDN is returned for local hosts.
NL_NUMERICHOST	If set, or if the host's name can't be located in the DNS, the numeric form

	of the host's address is returned instead of its name (e.g. by calling <i>inet_ntop()</i> instead of <i>getnodebyaddr()</i>).
NI_NAMEREQD	If set, an error is returned when the host's name can't be located in the DNS.
NI_NUMERICSERV	If set, the numeric form of the service address (instead of its name) is returned e.g. its port number. You may require two NI_NUMERICxxx flags to support the <i>-n</i> flag that many commands provide.
NI_DGRAM	Specify that the service is a datagram service. Call <i>getservbyport()</i> with a second argument of <i>udp</i> instead of its default of <i>tcp</i> . This is required for the few ports (512-514) that have different services for UDP and TCP.

Library:

`libsocket`

Description:

The *getnameinfo()* function defines and performs protocol-independent address-to-nodename translation. You can think of it as implementing the reverse-functionality of *getaddrinfo()* or similar functionality of *gethostbyaddr()* or *getservbyport()*.

This function looks up an IP address and port number provided by the caller in the DNS and system-specific database. For both IP address and port number, the *getnameinfo()* function returns text strings in respective buffers provided by the caller. The function indicates successful completion by a zero return value; a non-zero return value indicates failure.

The *getnameinfo()* function returns the nodename associated with the IP address in the buffer pointed to by the *host* argument. The *hostlen* argument gives the length of this buffer.

The *getnameinfo()* function returns the service name associated with the port number in the buffer pointed to by the *serv* argument. The *servlen* argument gives the length of this buffer.

Specify zero for *hostlen* or *servlen* when the caller chooses not to return either string. Otherwise, the caller must provide buffers large enough to hold the nodename and the service name, including the terminating null characters.

Most systems don't provide constants that specify the maximum size of either a FQDN or a service name. In order to aid your application in allocating buffers, the following constants are defined in `<netdb.h>`:

```
#define NI_MAXHOST  1025
#define NI_MAXSERV  32
```

You may find the first value as the constant `MAXDNAME` in recent versions of BIND's `<arpa/nameser.h>`; older versions of BIND define this constant to be 256. The second value is a guess based on the services listed in the current Assigned Numbers RFC. BIND (Berkeley Internet Name Domain) is a suite of functionalities that implements Domain Name System (DNS) protocols.

Extension

The implementation allows experimental numeric IPv6 address notation with scope identifier. An IPv6 link-local address appears as string like `fe80::1%ne0`, when the `NL_WITHSCOPEID` bit is enabled in the *flags* argument. See *getaddrinfo()* for the notation.

Returns:

0 Success.

Non-zero value

 An error occurred (see below).

Errors:

EALAGAIN	The name couldn't be resolved at this time. Future attempts may succeed.
EALBADFLAGS	The flags had invalid values.
EALFAIL	A nonrecoverable error occurred.
EALFAMILY	The address family wasn't recognized or the address length was invalid for the specified family.
EALMEMORY	There was a memory allocation failure.
EALNONAME	The name doesn't resolve for the supplied parameters. NI_NAMEREQD is set and the host's name can't be located, or both <i>node</i> name and <i>serv</i> name were null.
EALSYSTEM	A system error occurred. The error code can be found in <i>errno</i> .

Examples:

The following code gets the numeric hostname and the service name for a given socket address. There is no hardcoded reference to a particular address family.

```
struct sockaddr *sa; /* input */
char hbuf[NI_MAXHOST], sbuf[NI_MAXSERV];

if (getnameinfo(sa, sa->sa_len, hbuf, sizeof(hbuf), sbuf,
sizeof(sbuf), NI_NUMERICHOST | NI_NUMERICSERV) {
    errx(1, "could not get numeric hostname");
    /*NOTREACHED*/
}
printf("host=%s, serv=%s\n", hbuf, sbuf);
```

The following version checks if the socket address has reverse address mapping.

```
struct sockaddr *sa; /* input */
char hbuf[NI_MAXHOST];
```



```
if (getnameinfo(sa, sa->sa_len, hbuf, sizeof(hbuf), NULL, 0,
NI_NAMEREQD)) {
    errx(1, "could not resolve hostname"); /*NOTREACHED*/
}
printf("host=%s\n", hbuf);
```

Classification:

POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getaddrinfo(), *gethostbyaddr()*, *getservbyport()*, */etc/hosts*,
/etc/resolv.conf, */etc/services*, *named*

getnetbyaddr()

© 2004, QNX Software Systems Ltd.

Get a network entry, given an address (Unix)

Synopsis:

```
#include <netdb.h>

struct netent * getnetbyaddr( uint32_t net,
                             int type );
```

Arguments:

net The net address whose network entry you want to find.

type The address type. This must currently be AF_INET.

Library:

`libsocket`

Description:

The *getnetbyaddr()* function gets an entry for the given address, *net*, from the network database, `/etc/networks`.

This function returns a pointer to a structure of type `netent`, which contains the broken-out fields of a line in the network database.

The *setnetent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setnetent()*, the network database isn't closed after each call to *getnetbyname()*, or *getnetbyaddr()*.

The *getnetbyname()* and *getnetbyaddr()* functions sequentially search from the beginning of the file until a matching net name or net address and type is found, or until EOF is encountered. Network numbers are supplied in host order.

Returns:

A pointer to a valid `netent` structure, or NULL if an error occurs.

Files:

`/etc/networks`

Network name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

Only Internet network numbers are currently understood.

See also:

endnetent(), *getnetbyname()*, *getnetent()*, **netent**, *setnetent()*

`/etc/networks` in the *Utilities Reference*

getnetbyname()

© 2004, QNX Software Systems Ltd.

Get a network entry, given a name

Synopsis:

```
#include <netdb.h>

struct netent * getnetbyname( const char * name );
```

Arguments:

name The name of the network whose entry you want to find.

Library:

`libsocket`

Description:

The *getnetbyname()* function gets the network entry for the given name. This function returns a pointer to a structure of type `netent`, which contains the broken-out fields of a line in the network database, `/etc/networks`.

The *setnetent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setnetent()*, the network database isn't closed after each call to *getnetbyname()* or *getnetbyaddr()*.

The *getnetbyaddr()* and *getnetbyname()* functions sequentially search from the beginning of the file until a matching net name or net address and type is found, or until EOF is encountered. Network numbers are supplied in host order.

Returns:

A pointer to a valid `netent` structure, or NULL if an error occurs.

Files:

`/etc/networks`
Network name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

*endnetent(), getnetbyaddr(), getnetent(), **netent**, setnetent()*
/etc/networks in the *Utilities Reference*

getnetent()

© 2004, QNX Software Systems Ltd.

Read the next line of the network name database file

Synopsis:

```
#include <netdb.h>

struct netent * getnetent( void );
```

Library:

libsocket

Description:

The *getnetent()* function reads the next line of the network name database file, opening the file if necessary. It returns a pointer to a structure of type **netent**, which contains the broken-out fields of a line in the network database, */etc/networks*.

Returns:

A pointer to a valid **netent** structure, or NULL if an error occurs.

Files:

/etc/networks
Network name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

endnetent(), *getnetbyaddr()*, *getnetbyname()*, **netent**, *setnetent()*
/etc/networks in the *Utilities Reference*

getopt()

© 2004, QNX Software Systems Ltd.

Parse options from a command line

Synopsis:

```
#include <unistd.h>

int getopt( int argc,
            char * const argv[],
            const char * optstring );

extern char * optarg;
extern int optind, opterr, optopt;
```

Arguments:

argc The argument count that was passed to *main()*.

argv The argument array that was passed to *main()*.

optstring A string of recognized option letters; if a letter is followed by a colon, the option takes an argument. Valid option characters for *optstring* consist of a single alphanumeric character (i.e. a letter or digit).

Library:

`libc`

Description:

The *getopt()* function is a command-line parser that can be used by applications that follow the Utility Syntax Guidelines described below.

The *optind* global variable is the index of the next element of the *argv[]* vector to be processed. The system initializes *optind* to 1 when the program is loaded, and *getopt()* updates it when it finishes with each element of *argv[]*. Reset *optind* to 1 if you want to use *getopt()* to process additional argument sets.

The *getopt()* function returns the next option character from *argv* that matches a letter in *optstring*, if there's one that matches. If the option

takes an argument, *getopt()* sets the global variable *optarg* to point to the option argument as follows:

- 1 If the option is the last letter in the string pointed to by an element of *argv*, then *optarg* contains the next element of *argv*, and *optind* is incremented by 2.
- 2 Otherwise, *optarg* points to the string following the option letter in that element of *argv*, and *optind* is incremented by 1.

The *getopt()* function returns -1 is returned and doesn't change *optind* if:

- *argv[optind]* is NULL
- **argv[optind]* isn't the character '-'
- *argv[optind]* points to the string "--".

This function returns -1 after incrementing *optind*, if:

- *argv[optind]* points to the string "--".

If *getopt()* encounters an option character that isn't contained in *optstring*, it returns the ? character. If it detects a missing option argument, it returns : if the first character of *optstring* is a colon, or ? otherwise. In both cases, *getopt()* sets *optopt* to the option character that caused the error.

The *getopt()* always prints a diagnostic message to *stderr* unless *opterr* is set to 0, or the first character of *optstring* is a : character.

Utility Syntax Guidelines

The *getopt()* function may be used by applications that follow these guidelines:

- When describing the syntax of a utility, the options are listed in alphabetical order. There's no implied relationship between the options based upon the order in which they appear, unless otherwise stated in the Options section, or:

- the options are documented as mutually-exclusive and such an option is documented to override any incompatible options preceding it
- when an option has option arguments repeated, the option and option argument combinations are interpreted in the order specified on the command line.

If an option that doesn't have option arguments is repeated, the results depend on the application.

- Names of parameters that require substitution by actual values may be shown with embedded underscores or as *<parameter name>*. Angle brackets are used for the symbolic grouping of a phrase representing a single parameter and portable applications shouldn't include them in data submitted to the utility.
- Options may be documented individually, or grouped (if they don't take option arguments):

```
utility_name [-a] [-b] [-c option_argument]
              [-d|-e] [-foption_argument] [operand...]
```

Or:

```
utility_name [-ab] [-c option_argument]
              [-d|-e] [-foption_argument] [operand...]
```

Utilities with very complex arguments may be shown as:

```
utility_name [options] [operand]
```

- Unless specified, whenever an operand or option argument is, or contains, a numeric value:
 - the number is interpreted as a decimal integer
 - numerals in the range 0 to 2,147,483,647 are syntactically recognized as numeric values
 - when the utility description states that it accepts negative numbers as operands or option arguments, numerals in the

range -2,147,483,647 to 2,147,483,647 are syntactically recognized as numeric values

- ranges greater than those listed here are allowed.

All numbers within the allowable range aren't necessarily semantically correct. A standard utility that accepts an option argument or operand that's to be interpreted as a number, and for which a range of values smaller than that shown above is permitted, describes that smaller range along with the description of the option argument or operand. If an error is generated, the utility's diagnostic message indicates that the value is out of the supported range, not that it's syntactically incorrect.

- Arguments or option arguments enclosed in the “[” and “]” notation are optional and can be omitted. Portable applications shouldn't include the “[” and “]” symbols in data submitted to the utility.
- Ellipses (...) are used to denote that one or more occurrences of an option or operand are allowed. When an option or an operand followed by ellipses is enclosed in brackets, zero or more options or operands may be specified. The forms:

```
utility_name -f option_argument...[operand...]  
utility_name [-g option_argument]...[operand...]
```

indicate that multiple occurrences of the option and its option argument preceding the ellipses are valid, with semantics as indicated in the Options section of the utility. In the first example, each option argument requires a preceding -f and at least one -foption_argument must be given.

- When the synopsis is too long to be printed on a single line in the documentation, the indented lines following the initial line are continuation lines. An actual use of the command appears on a single logical line.

Returns:

The next option character specified on the command line; a colon if a missing argument is detected and the first character of *optstring* is a colon; a question mark if an option character is encountered that's not in *optstring* and the first character of *optstring* isn't a colon; otherwise, -1 when all command line options have been parsed.

Examples:

```
#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>

int main( int argc, char* argv[] )
{
    int c, errflag = 0;

    while( ( c = getopt( argc, argv, "abt:" ) )
        != -1 ) {
        switch( c ) {
            case 'a': printf( "apples\n" );
                     break;
            case 'b': printf( "bananas\n" );
                     break;
            case 't': printf( "tree = %s\n", optarg );
                     break;
            case '?': ++errflag;
                     break;
        }
    }
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1a

Safety

Cancellation point Yes

Interrupt handler No

continued...

Safety

Signal handler	No
Thread	No

See also:

getsubopt(), stderr

Guidelines 3,4,5,6,7,9 and 10 in the *Base Definitions* volume of the IEEE Std.1003.1-2001, Section 12.2, *Utility Syntax Guidelines*.

getpass()

© 2004, QNX Software Systems Ltd.

Prompt for and read a password

Synopsis:

```
#include <unistd.h>

char *getpass( const char *prompt );
```

Arguments:

prompt The string you want to display to prompt for the password.

Library:

`libc`

Description:

The *getpass()* function can be used to get a password. It opens the current terminal, displays the given *prompt*, suppresses echoing, reads up to 32 characters into a static buffer, and restores echoing. This function adds a null character to the end of the string, but ignores additional characters and the newline character.

Returns:

A pointer to the static buffer.

Classification:

Legacy Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function leaves its result in an internal static buffer and returns a pointer to it. Subsequent calls to *getpass()* modify the same buffer. The calling process should zero the password as soon as possible to avoid leaving the clear-text password visible in the process's address space.

See also:

crypt()

getpeername()

© 2004, QNX Software Systems Ltd.

Get the name of the peer connected to a socket

Synopsis:

```
#include <sys/socket.h>

int getpeername( int s,
                 struct sockaddr * name,
                 socklen_t * namelen );
```

Arguments:

<i>s</i>	The socket whose connected peer you want to get.
<i>name</i>	A buffer where the function can store the name of the peer.
<i>namelen</i>	A pointer to a <code>socklen_t</code> object that initially specifies the size of the buffer. This function stores the actual size of the name, in bytes, in this object.

Library:

`libsocket`

Description:

The `getpeername()` function returns the name of the peer connected to socket *s*. The name is truncated if the buffer provided is too small.

Returns:

0	Success.
-1	An error occurred (<i>errno</i> is set).

Errors:

EBADF	Invalid descriptor <i>s</i> .
EFAULT	The <i>name</i> parameter points to memory <i>not</i> in a valid part of the process address space.

ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOTCONN	The socket isn't connected.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

accept(), bind(), getsockname(), socket()

getpgid()

© 2004, QNX Software Systems Ltd.

Get a process group ID

Synopsis:

```
#include <unistd.h>

pid_t getpgid( pid_t pid );
```

Arguments:

pid The ID of the process whose process group ID you want to get.

Library:

`libc`

Description:

The *getpgid()* returns the group ID for the process specified by *pid*. If *pid* is 0, *getpgid()* returns the calling process's group ID.

The following definitions are worth mentioning:

Process	An executing instance of a program, identified by a nonnegative integer called a process ID.
Process group	A collection of one or more processes, with a unique process group ID. A process group ID is a positive integer.

Returns:

A process group ID for success, or `(pid_t)-1` if an error occurs.

Errors:

If an error occurs, *errno* is set to:

ESRCH The process specified by *pid* doesn't exist.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

getsid(), *setpgid()*, *setsid()*

getpgrp()

© 2004, QNX Software Systems Ltd.

Get the process group

Synopsis:

```
#include <sys/types.h>
#include <process.h>

pid_t getpgrp( void );
```

Library:

libc

Description:

The *getpgrp()* function gets the ID of the process group to which the calling process belongs.

Returns:

The calling process's process group ID.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <process.h>
#include <sys/types.h>

int main( void )
{
    printf( "I am in process group %d\n", (int) getpgrp() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point No

continued...

Safety

Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

setpgrp(), *setsid()*

getpid()

Get the process ID

© 2004, QNX Software Systems Ltd.

Synopsis:

```
#include <process.h>

pid_t getpid( void );
```

Library:

libc

Description:

The *getpid()* function gets the process ID for the calling process.

Returns:

The process ID of the calling process.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <process.h>

int main ( void )
{
    printf( "I'm process %d\n", getpid() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes

continued...

Safety

Thread	Yes
--------	-----

See also:

getppid()

getppid()

© 2004, QNX Software Systems Ltd.

Get the parent process ID

Synopsis:

```
#include <sys/types.h>
#include <process.h>

pid_t getppid( void );
```

Library:

libc

Description:

The *getppid()* function gets the process ID of the parent of the calling process.

Returns:

The calling process's parent's process ID.

Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <process.h>

int main( void )
{
    printf( "My parent is %d\n", getppid() );
    return EXIT_SUCCESS;
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point No

continued...

Safety

Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

getpid()

getprio()

© 2004, QNX Software Systems Ltd.

Get the priority of a given process

Synopsis:

```
#include <sched.h>

int getprio( pid_t pid );
```

Arguments:

pid The process ID of the process whose priority you want to get.

Library:

`libc`

Description:

The *getprio()* function returns the current priority of thread 1 in process *pid*. If *pid* is zero, the priority of the calling thread is returned.

Returns:

The priority, or -1 if an error occurred (*errno* is set).

Errors:

ESRCH The process *pid* doesn't exist.

Classification:

QNX 4

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

The *getprio()* and *setprio()* functions are included in the QNX Neutrino libraries for porting QNX 4 applications. For new programs, use *sched_getparam()* or *pthread_getschedparam()*.

See also:

errno, *pthread_getschedparam()*, *pthread_setschedparam()*, *sched_get_priority_max()*, *sched_get_priority_min()*, *sched_getparam()*, *sched_getscheduler()*, *sched_setscheduler()*, *sched_yield()*, *setprio()*

getprotobyname()

© 2004, QNX Software Systems Ltd.

Get a protocol entry, given a name

Synopsis:

```
#include <netdb.h>

struct protoent * getprotobyname( const char * name );
```

Arguments:

name The name of the protocol whose entry you want to get.

Library:

libsocket

Description:

The *getprotobyname()* function gets the entry for the given name from the protocol database, */etc/protocols*. This function returns a pointer to a structure of type **protoent**, which contains the broken-out fields of a line in the network protocol database.

The *setprotoent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setprotoent()*, the protocol database isn't closed after each call to *getprotobyname()* or *getprotobynumber()*.

The *getprotobyname()* and *getprotobynumber()* functions sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

Returns:

A pointer to a valid **protoent** structure, or NULL if an error occurs.

Files:

/etc/protocols
Protocol name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

Currently, only the Internet protocols are understood.

See also:

endprotoent(), *getprotobynumber()*, *getprotoent()*, **protoent**,
setprotoent()

/etc/protocols in the *Utilities Reference*

getprotobynumber()

© 2004, QNX Software Systems Ltd.

Get a protocol entry, given a number

Synopsis:

```
#include <netdb.h>

struct protoent * getprotobynumber( int proto );
```

Arguments:

proto The protocol number whose entry you want to get.

Library:

`libsocket`

Description:

The *getprotobynumber()* function gets the protocol entry for the given number. It returns a pointer to structure of type `protoent`, which contains the broken-out fields of a line in the network protocol database, `/etc/protocols`.

The *setprotoent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setprotoent()*, the protocol database isn't closed after each call to *getprotobyname()* or *getprotobynumber()*.

The *getprotobyname()* and *getprotobynumber()* functions sequentially search from the beginning of the file until a matching protocol name or protocol number is found, or until EOF is encountered.

Returns:

A pointer to a valid `protoent` structure, or NULL if an error occurs.

Files:

`/etc/protocols`
Protocol name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

Currently, only the Internet protocols are understood.

See also:

endprotoent(), *getprotobyname()*, *getprotoent()*, **protoent**, *setprotoent()*

/etc/protocols in the *Utilities Reference*

getprotoent()

© 2004, QNX Software Systems Ltd.

Read the next line of the protocol name database file

Synopsis:

```
#include <netdb.h>

struct protoent * getprotoent( void );
```

Library:

libsocket

Description:

The *getprotoent()* function reads the next line of the protocol name database file, opening the file if necessary. It returns a pointer to a structure of type **protoent**, which contains the broken-out fields of a line in the network protocol database, */etc/protocols*.

Returns:

A pointer to a valid **protoent** structure, or NULL if an error occurs.

Files:

/etc/protocols
Protocol name database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

Currently, only the Internet protocols are understood.

See also:

endprotoent(), *getprotobyname()*, *getprotobynumber()*, **protoent**, *setprotoent()*

/etc/protocols in the *Utilities Reference*

getpwent()

© 2004, QNX Software Systems Ltd.

Get an entry from the password database

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

struct passwd* getpwent( void );
```

Library:

libc

Description:

The *getpwent()* function returns the next entry from the password database. This function uses a static buffer that's overwritten by each call.



The *getpwent()*, *getpwnam()*, and *getpwuid()*, functions share the same static buffer.

Returns:

A pointer to an object of type `struct passwd` containing the next entry from the password database. When *getpwent()* is first called, the password database is opened, and remains open until either a NULL is returned to signify end-of-file, or *endpwent()* is called.

Errors:

The *getpwent()* function uses the following functions, and as a result, *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

Examples:

```

/*
 * This program loops, reading a login name from standard
 * input and checking to see if it is a valid name. If it
 * is not valid, the entire contents of the name in the
 * password database are printed.
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <pwd.h>

int main( void )
{
    struct passwd* pw;
    char    buf[80];

    setpwent( );
    while( gets( buf ) != NULL ) {
        if( ( pw = getpwnam( buf ) ) != ( struct passwd * )0 ) {
            printf( "Valid login name is: %s\n", pw->pw_name );
        } else {
            setpwent( );
            while( ( pw=getpwent( ) ) != ( struct passwd * )0 )
                printf( "%s\n", pw->pw_name );
        }
    }
    endpwent();
    return( EXIT_SUCCESS );
}

```

Classification:

Standard Unix

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

endpwent(), *errno*, *getgrent()*, *getlogin()*, *getpwnam()*, *getpwuid()*,
setpwent()

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

struct passwd* getpwnam( const char* name );
```

Arguments:

name The name of the user whose entry you want to find.

Library:

`libc`

Description:

The *getpwnam()* function gets information about the user with the given *name*. It uses a static buffer that's overwritten by each call.



The *getpwent()*, *getpwnam()*, and *getpwuid()* functions share the same static buffer.

The *getpwnam_r()* function is a reentrant version of *getpwnam()*.

Returns:

A pointer to an object of type `struct passwd` containing an entry from the group database with a matching *name*. A NULL pointer is returned on error or failure to find a entry with a matching *name*.

Examples:

```
/*
 * Print information from the password entry
 * about the user name given as argv[1].
 */
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>
#include <sys/types.h>
```

```
#include <pwd.h>

int main( int argc, char* *argv )
{
    struct passwd* pw;

    if( ( pw = getpwnam( argv[1] ) ) == NULL ) {
        fprintf( stderr, "getpwnam: unknown %s\n",
            argv[1] );
        return( EXIT_FAILURE );
    }
    printf( "login name  %s\n", pw->pw_name );
    printf( "user id     %d\n", pw->pw_uid );
    printf( "group id    %d\n", pw->pw_gid );
    printf( "home dir   %s\n", pw->pw_dir );
    printf( "login shell %s\n", pw->pw_shell );
    return( EXIT_SUCCESS );
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getlogin(), getpwent(), getpwnam_r() getpwuid()

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

int getpwnam_r( const char* name,
                struct passwd* pwd,
                char* buffer,
                size_t bufsize,
                struct passwd* result );
```

Arguments:

- name* The name of the user whose entry you want to find.
- pwd* A pointer to a **passwd** structure where the function can store the entry.
- buffer* A block of memory that the function can use to allocate storage referenced by the **passwd** structure. You can determine the maximum size needed for this buffer by calling *sysconf()* with an argument of **_SC_GETPW_R_SIZE_MAX**.
- bufsize* The size of the block that *buffer* points to, in characters.
- result* The address of a pointer to a **passwd** structure. If *getpwnam_r()* finds the entry, it stores a pointer to *pwd* in the location indicated by *result*; otherwise the function stores a NULL pointer there.

Library:

libc

Description:

The *getpwnam_r()* function is a reentrant version of *getpwnam()*. It gets information about the user with the given *name*.

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, the `getpwnam_r()` function updates the `passwd` structure pointed to by `pwd` and stores a pointer to that structure at the location pointed by `result`. The structure contains an entry from the user database with the given `name`.

The function stores a NULL pointer at the location pointed by `result` on error or if it can't find the requested entry.

Returns:

Zero for success, or an error number.

Errors:

ERANGE Insufficient storage was supplied via `buffer` and `bufsize` to contain the resulting `passwd` structure.

The `getpwnam_r()` function uses the following functions, and as a result, `errno` can be set to an error for any of these calls:

- `fclose()`
- `fgets()`
- `fopen()`
- `fseek()`
- `rewind()`

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getlogin(), getpwent(), getpwnam(), getpwuid(), getpwuid_r()

getpwuid()

© 2004, QNX Software Systems Ltd.

Get information about the user with a given ID

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

struct passwd* getpwuid( uid_t uid );
```

Arguments:

uid The userid whose entry you want to find.

Library:

`libc`

Description:

The *getpwuid()* function gets information about user *uid*. This function uses a static buffer that's overwritten by each call.



The *getpwent()*, *getpwnam()*, and *getpwuid()* functions share the same static buffer.

Returns:

A pointer to an object of type `struct passwd` containing an entry from the group database with a matching *uid*, or NULL if an error occurred or the function couldn't find a matching entry.

Examples:

```
/*
 * Print password info on the current user.
 */
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>
#include <sys/types.h>
#include <pwd.h>

int main( void )
```

```
{
    struct passwd* pw;

    if( ( pw = getpwuid( getuid() ) ) == NULL ) {
        fprintf( stderr,
            "getpwuid: no password entry\n" );
        return( EXIT_FAILURE );
    }
    printf( "login name  %s\n", pw->pw_name );
    printf( "user id     %d\n", pw->pw_uid );
    printf( "group id     %d\n", pw->pw_gid );
    printf( "home dir     %s\n", pw->pw_dir );
    printf( "login shell  %s\n", pw->pw_shell );
    return( EXIT_SUCCESS );
}
```

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

See also:

getlogin(), getpwent(), getpwnam()

getpwuid_r()

© 2004, QNX Software Systems Ltd.

Get information about the user with a given ID

Synopsis:

```
#include <sys/types.h>
#include <pwd.h>

int getpwuid_r( uid_t uid,
               struct passwd* pwd,
               char* buffer,
               size_t bufsize,
               struct passwd** result );
```

Arguments:

<i>uid</i>	The userid whose entry you want to find.
<i>pwd</i>	A pointer to a passwd structure where the function can store the entry.
<i>buffer</i>	A block of memory that the function can use to allocate storage referenced by the passwd structure. You can determine the maximum size needed for this buffer by calling <i>sysconf()</i> with an argument of <code>_SC_GETPW_R_SIZE_MAX</code> .
<i>bufsize</i>	The size of the block that <i>buffer</i> points to, in characters.
<i>result</i>	The address of a pointer to a passwd structure. If <i>getpwnam_r()</i> finds the entry, it stores a pointer to <i>pwd</i> in the location indicated by <i>result</i> ; otherwise the function stores a NULL pointer there.

Library:

libc

Description:

The *getpwuid_r()* function is a reentrant version of *getpwuid()*. It lets a process gain more knowledge about user with the given *uid*.

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, the `getpwuid_r()` function updates the `passwd` structure pointed to by `pwd` and stores a pointer to that structure at the location pointed by `result`. The structure contains an entry from the user database with a matching `uid`.

The function stores a NULL pointer at the location pointed by `result` on error or if it can't find the requested entry.

Returns:

Zero for success, or an error number.

Errors:

`ERANGE` Insufficient storage was supplied via `buffer` and `bufsize` to contain the resulting `passwd` structure.

The `getpwuid_r()` function uses the following functions, and as a result, `errno` can be set to an error for any of these calls:

- `fclose()`
- `fgets()`
- `fopen()`
- `fseek()`
- `rewind()`

Classification:

POSIX 1003.1

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getlogin(), getpwnam(), getpwnam_r(), getpwuid()

Synopsis:

```
#include <sys/resource.h>

int getrlimit( int resource,
              struct rlimit * rlp );

int getrlimit64( int resource,
                struct rlimit64 * rlp );
```

Arguments:

resource The resource whose limit you want to get. For a list of the possible resources, their descriptions, and the actions taken when the current limit is exceeded, see *setrlimit()*.

rlp A pointer to a `rlimit` or `rlimit64` structure where the function can store the limit on the resource. The `rlimit` and `rlimit64` structures include at least the following members:

```
rlim_t rlim_cur; /* current (soft) limit */
rlim_t rlim_max; /* hard limit */
```

The `rlim_t` type is an arithmetic data type to which you can cast objects of type `int`, `size_t`, and `off_t` without loss of information.

Library:

`libc`

Description:

The *getrlimit()* function gets the limits on the consumption of a variety of system resources by a process and each process it creates. The *getrlimit64()* function is a 64-bit version of *getrlimit()*.

Each call to *getrlimit()* identifies a specific resource to be operated upon as well as a resource limit. A resource limit is a pair of values:

- the current (soft) limit
- a maximum (hard) limit.

A process can change soft limits to any value that's less than or equal to the hard limit. A process may (irreversibly) lower its hard limit to any value that's greater than or equal to the soft limit. Only a process with an effective user ID of `root` can raise a hard limit. Both hard and soft limits can be changed in a single call to *setrlimit()* subject to the constraints described above. Limits may have an "infinite" value of `RLIM_INFINITY`.

Because limit information is stored in the per-process information, the shell builtin `ulimit` command must directly execute this system call if it's to affect all future processes created by the shell.

The values of the current limit of the following resources affect these parameters:

Resource	Parameter
<code>RLIMIT_FSIZE</code>	<code>FCHR_MAX</code>
<code>RLIMIT_NOFILE</code>	<code>OPEN_MAX</code>

When using *getrlimit()*, if a resource limit can be represented correctly in an object of type `rlim_t`, then its representation is returned; otherwise, if the value of the resource limit is equal to that of the corresponding saved hard limit, the value returned is `RLIM_SAVED_MAX`; otherwise, the value returned is `RLIM_SAVED_CUR`.

A limit whose value is greater than `RLIM_INFINITY` is permitted.

The *exec** family of functions also causes resource limits to be saved.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EFAULT The *rlp* argument points to an illegal address.
- EINVAL An invalid resource was specified.
- EPERM The limit specified to *setrlimit()* would've raised the maximum limit value, and the effective user of the calling process isn't the superuser.

Classification:

getrlimit() is standard Unix; *getrlimit64()* is for large-file support

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

See also:

brk(), *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *fork()*, *getdtablesize()*, *malloc()*, *open()*, *setrlimit()*, *setrlimit64()*, *signal()*, *sysconf()*

getrusage()

© 2004, QNX Software Systems Ltd.

Get information about resource utilization

Synopsis:

```
#include <sys/resource.h>

int getrusage( int who,
               struct rusage * r_usage );
```

Arguments:

who Which process to get the usage for:

- RUSAGE_CHILDREN — get information about resources used by the terminated and waited-for children of the current process. If the child is never waited for (e.g if the parent has SA_NOCLDWAIT set, or sets SIGCHLD to SIG_IGN), the resource information for the child process is discarded and isn't included.
- RUSAGE_SELF — get information about resources used by the current process.

r_usage A pointer to an object of type **struct rusage** in which the function can store the resource information; see below.

Library:

libc

Description:

The *getrusage()* function provides measures of the resources used by the current process or its terminated and waited-for child processes, depending on the value of the *who* argument.

The **rusage** structure is defined as:

```
struct timeval ru_utime; /* user time used */
struct timeval ru_stime; /* system time used */
long          ru_maxrss; /* max resident set size */
```

```

long      ru_ixrss;    /* integral shared memory size */
long      ru_idrss;    /* integral unshared data " */
long      ru_isrss;    /* integral unshared stack " */
long      ru_minflt;   /* page reclaims */
long      ru_majflt;   /* page faults */
long      ru_nswap;    /* swaps */
long      ru_inblock;  /* block input operations */
long      ru_oublock;  /* block output operations */
long      ru_msgsnd;   /* messages sent */
long      ru_msgrcv;   /* messages received */
long      ru_nsignals; /* signals received */
long      ru_nvcsw;    /* voluntary context switches */
long      ru_nivcsw;   /* involuntary " */

```

The members include:

<i>ru_utime</i>	The total amount of time, in seconds and microseconds, spent executing in user mode.
<i>ru_stime</i>	The total amount of time, in seconds and microseconds, spent executing in system mode.
<i>ru_maxrss</i>	The maximum resident set size, given in pages. See the Caveats section, below.
<i>ru_ixrss</i>	Not currently supported.
<i>ru_idrss</i>	An “integral” value indicating the amount of memory in use by a process while the process is running. This value is the sum of the resident set sizes of the process running when a clock tick occurs. The value is given in pages times clock ticks. It doesn’t take sharing into account. See the Caveats section, below.
<i>ru_isrss</i>	Not currently supported.
<i>ru_minflt</i>	The number of page faults serviced that didn’t require any physical I/O activity. See the Caveats section, below.
<i>ru_majflt</i>	The number of page faults serviced that required physical I/O activity. This could include page ahead

operations by the kernel. See the Caveats section, below

<i>ru_nswap</i>	The number of times a process was swapped out of main memory.
<i>ru_inblock</i>	The number of times the file system had to perform input in servicing a <i>read()</i> request.
<i>ru_oublock</i>	The number of times the filesystem had to perform output in servicing a <i>write()</i> request.
<i>ru_msgsnd</i>	The number of messages sent over sockets.
<i>ru_msgrcv</i>	The number of messages received from sockets.
<i>ru_nsignals</i>	The number of signals delivered.
<i>ru_nvcsw</i>	The number of times a context switch resulted due to a process's voluntarily giving up the processor before its timeslice was completed (usually to await availability of a resource).
<i>ru_nivcsw</i>	The number of times a context switch resulted due to a higher priority process's becoming runnable or because the current process exceeded its time slice.

Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

Errors:

- EFAULT The address specified by the *r_usage* argument isn't in a valid portion of the process's address space.
- EINVAL Invalid *who* parameter.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	Yes
Thread	Yes

Caveats:

Only the `timeval` fields of `struct rusage` are supported.

The numbers `ru_inblock` and `ru_oublock` account only for real I/O, and are approximate measures at best. Data supplied by the cache mechanism is charged only to the first process to read and the last process to write the data.

The way resident set size is calculated is an approximation, and could misrepresent the true resident set size.

Page faults can be generated from a variety of sources and for a variety of reasons. The customary cause for a page fault is a direct reference by the program to a page that isn't in memory. Now, however, the kernel can generate page faults on behalf of the user, for example, servicing `read()` and `write()` functions. Also, a page fault can be caused by an absent hardware translation to a page, even though the page is in physical memory.

In addition to hardware-detected page faults, the kernel may cause pseudo page faults in order to perform some housekeeping. For example, the kernel may generate page faults, even if the pages exist in physical memory, in order to lock down pages involved in a raw I/O request.

By definition, major page faults require physical I/O, while minor page faults don't. For example, reclaiming the page from the free list would avoid I/O and generate a minor page fault. More commonly, minor page faults occur during process startup as references to pages which are already in memory. For example, if an address space faults on some "hot" executable or shared library, a minor page fault results for the address space. Also, anyone doing a *read()* or *write()* to something that's in the page cache gets a minor page fault(s) as well.

There's no way to obtain information about a child process that hasn't yet terminated.

See also:

gettimeofday(), *read()*, *times()*, *wait()*, *write()*

Synopsis:

```
#include <stdio.h>

char *gets( char *buf );
```

Arguments:

buf A buffer where the function can store the string.

Library:

`libc`

Description:

The *gets()* function gets a string of characters from the *stdin* stream, and stores them in the array pointed to by *buf* until end-of-file is encountered or a newline character is read. Any newline character is discarded, and the string is NUL-terminated.



You should use *fgets()* instead of *gets()*; *gets()* happily overflows the *buf* array if a newline character isn't read from *stdin* before the end of the array is reached.

The *gets()* function is similar to *fgets()*, except that *gets()* operates with *stdin*, has no size argument, and replaces a newline character with the NUL character.

Returns:

A pointer to *buf*, or NULL when end-of-file is encountered before reading any characters or a read error occurred (*errno* is set).



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main( void )
{
    char buffer[80];

    while( gets( buffer ) != NULL ) {
        puts( buffer );
    }

    return EXIT_SUCCESS;
}
```

Classification:

ANSI

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, *feof()*, *ferror()*, *fopen()*, *getc()*, *fgetc()*, *fgets()*, *puts()*, *ungetc()*

Synopsis:

```
#include <netdb.h>

struct servent * getservbyname( const char * name,
                                const char * proto );
```

Arguments:

name The name of the service whose entry you want to find.

proto NULL, or the protocol for the service.

Library:

libsocket

Description:

The *getservbyname()* function gets the entry for the given name and protocol from the network services database, */etc/services*. This function returns a pointer of type **servent**, which contains the broken-out fields of a line in the network services database.

The *setservent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setservent()*, the services database isn't closed after each call to *getservbyname()* or *getservbyport()*.

The *getservbyname()* and *getservbyport()* functions sequentially search from the beginning of the file until a matching protocol name or port number is found, or until EOF is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

Returns:

A valid pointer to a **servent** structure, or NULL if an error occurs.

Files:

`/etc/services`

Network services database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point Yes

Interrupt handler No

Signal handler No

Thread No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

See also:

endservent(), *getprotoent()*, *getservbyport()*, *getservent()*, **servent**, *setservent()*

`/etc/services` in the *Utilities Reference*

Synopsis:

```
#include <netdb.h>

struct servent * getservbyport( int port,
                               const char * proto );
```

Arguments:

port The port number for the service.

proto NULL, or the protocol for the service.

Library:

`libsocket`

Description:

The *getservbyport()* function gets the entry for the given port from the services database, `/etc/services`. This function returns a pointer to a structure of type `servent`, which contains the broken-out fields of a line in the network services database.

The *setservent()* function opens and rewinds the file. If you pass a nonzero *stayopen* argument to *setservent()*, the services database isn't closed after each call to *getservbyname()* or *getservbyport()*.

The *getservbyport()* function sequentially searches from the beginning of the file until a matching protocol name or port number is found, or until EOF is encountered. If a protocol name is also supplied (non-NULL), searches must also match the protocol.

Returns:

A valid pointer to a `servent` structure, or NULL if an error occurs.

Files:

`/etc/services`
Network services database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

See also:

endservent(), *getservbyname()*, *getservent()*, **servent**, *setservent()*
`/etc/services` in the *Utilities Reference*

Synopsis:

```
#include <netdb.h>

struct servent * getservent( void );
```

Library:

libsocket

Description:

The *getservent()* function reads the next line of network services database file, opening the file if necessary. It returns a pointer to a structure of type **servent**, which contains the broken-out fields of a line in the network services database, */etc/services*.

Returns:

A valid pointer to a **servent** structure, or NULL if an error occurs.

Files:

/etc/services
Network services database file.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	No

Caveats:

This function uses static data; if you need the data for future use, copy it before any subsequent calls overwrite it.

See also:

endservent(), *getservbyname()*, *getservbyport()*, **servent**,
setservent()

/etc/services in the *Utilities Reference*

Synopsis:

```
#include <unistd.h>

pid_t getsid( pid_t pid );
```

Arguments:

pid The process ID for the process whose session ID you want to get.

Library:

libc

Description:

The *getsid()* function determines the session ID for the given process ID, *pid*.

Returns:

The session ID, or -1 if an error occurs (*errno* is set).

Errors:

EPERM The process specified by *pid* is not in the same session as the calling process. The implementation doesn't allow access to the process group ID of the session leader from the calling process.

EINVAL There isn't a process with the given ID.

Classification:

Standard Unix

Safety

Cancellation point	No
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

errno, setsid()

Synopsis:

```
#include <sys/socket.h>

int getsockname( int s,
                 struct sockaddr * name,
                 socklen_t * namelen );
```

Arguments:

<i>s</i>	The file descriptor of the socket whose name you want to get.
<i>name</i>	A pointer to a sockaddr object where the function can store the socket's name.
<i>namelen</i>	A pointer to a socklen_t object that initially indicates the amount of space pointed to by <i>name</i> . The function updates <i>namelen</i> to contain the actual size of the name (in bytes).

Library:

libsocket

Description:

The *getsockname()* function returns the current name for the specified socket.

Returns:

0	Success.
-1	An error occurred (<i>errno</i> is set).

Errors:

EBADF	Invalid descriptor <i>s</i> .
EFAULT	The <i>name</i> parameter points to memory that isn't in a valid part of the process address space.
ENOBUFS	Insufficient resources were available in the system to perform the operation.

Classification:

Standard Unix, POSIX 1003.1-2001

Safety

Cancellation point	Yes
Interrupt handler	No
Signal handler	No
Thread	Yes

See also:

getpeername()

Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>

int getsockopt( int s,
               int level,
               int optname,
               void * optval,
               socklen_t * optlen );
```

Arguments:

- | | |
|----------------|--|
| <i>s</i> | The file descriptor of the socket that the option is to be applied to, as returned by <i>socket()</i> . |
| <i>level</i> | The protocol layer that the option is to be applied to. In most cases, it's a socket-level option and is indicated by SOL_SOCKET. |
| <i>optname</i> | The option for the socket file descriptor. For a list of options, see "Options," below. |
| <i>optval</i> | A pointer to the value of the option (in most cases, whether the option is to be turned on or off). If no option value is to be returned, <i>optval</i> may be NULL.

Most socket-level options use an <code>int</code> parameter for <i>optval</i> . Others, such as the SO_LINGER, SO_SNDTIMEO, and SO_RCVTIMEO options, use structures that also let you get data associated with the option. |
| <i>optlen</i> | A pointer to the length of the value of the option. This argument is a value-result parameter; initialize it to indicate the size of the buffer pointed to by <i>optval</i> . |

Library:

libsocket

Description:

The *getsockopt()* function gets options associated with a socket.

Manipulating socket options

When manipulating a socket option, you must specify the option's name (*optname*) and the level (*level*) at which the option resides.

To manipulate options at the socket-level, specify *level* as SOL_SOCKET. When manipulating options any other level, the value that you specify for *level* is represented by the protocol number of the appropriate protocol controlling the option. You can obtain the value in a variety of ways:

- from the “Options” section below, use the symbolic constant (e.g. IPPROTO_IP, IPPROTO_TCP) that corresponds to the option
- from `/etc/protocols`, specify the protocol number for the appropriate protocol
- call *getprotobyname()* and pass the appropriate protocol (e.g. `getprotobyname(tcp);`) to retrieve the number of the protocol level.



The latter two ways might not work if you have customized `/etc/protocols`.

The *optname* parameter and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The `<sys/socket.h>` header file contains definitions for the socket-level options. Options at other protocol levels vary in format and name.



Since levels (e.g. SOL_SOCKET, IPPROTO_IP and IPPROTO_TCP) and the options within the levels can vary, you need to ensure the proper headers are included for both. For example, when setting TCP_NODELAY:

```
int on = 1;
setsockopt(s, IPPROTO_TCP, TCP_NODELAY, &on);
```

the level IPPROTO_TCP is defined in `<netinet/in.h>`, whereas the TCP_NODELAY option is defined in `<netinet/tcp.h>`.

Options

Here are some of the more common options and their corresponding *level*.



Except where noted, you can examine the state of the option by calling *getsockopt()*, and set the state by calling *setsockopt()*.

For the list of options that the tiny TCP/IP stack supports, see `npm-ttcpip.so` in the *Utilities Reference*.

IP_HDRINCL

level: IPPROTO_IP

Get or set the custom IP header that's included with your data. You can use it only for raw sockets. For example:

```
(socket(AF_INET, SOCK_RAW, ...))
```

IP_TOS

level: IPPROTO_IP

Get or set the type-of-service field in the IP header for SOCK_STREAM and SOCK_DGRAM (not applicable for `npm-ttcpip.so`) sockets.

- SO_BINDTODEVICE** *level: SOL_SOCKET*
Applies to *setsockopt()* only.
- Allow packets to be sent or received on this specified interface only. If the interface specified conflicts with the parameters of *bind()*, or with the routing table, an error or undesired behavior may occur.
- This option accepts the **ifreq** structure with the *ifr_name* member set to the interface name (e.g. **en0**). Currently, you can use this option only for UDP sockets.
- SO_BROADCAST** *level: SOL_SOCKET*
Enable or disable the permission to transmit broadcast messages. You can use this option only for UDP sockets. For example:
- ```
(socket(AF_INET, SOCK_DGRAM, ...))
```
- “Broadcast” was a privileged operation in earlier versions of the system.
- SO\_DEBUG** *level: SOL\_SOCKET*  
Enable or disable the recording of debug information in the underlying protocol modules.
- SO\_DONTROUTE** *level: SOL\_SOCKET*  
Enable or disable the bypassing of routing tables for outgoing messages. Indicates that outgoing messages should bypass the standard routing facilities. The messages are directed to the appropriate network interface according to the network portion of the destination address.
- SO\_ERROR** *level: SOL\_SOCKET*  
Applies to *getsockopt()* only.
- Get any pending error on the socket and clears the error status. You can use it to check for asynchronous errors on connected datagram sockets or for other asynchronous errors.

- SO\_KEEPAIVE** *level: SOL\_SOCKET*
- Enable or disable the periodic (at least every 2 hours) transmission of messages on a connected socket. Should the connected party fail to respond to these messages, the connection is considered broken, and processes that are using the socket are notified via a SIGPIPE signal when they attempt to send data.
- SO\_LINGER** *level: SOL\_SOCKET*
- Controls the action that's taken when unsent messages are queued on socket when a *close()* is performed.
- If it's enabled and the socket promises reliable delivery of data, the system blocks the process on the *close()* attempt until it's able to transmit the data or until it decides it can't deliver the information (a timeout period, termed the linger interval, is specified in the *setsockopt()* call when SO\_LINGER is requested).
- If it's disabled, the system processes the *close()* in a way that lets the process continue as quickly as possible.
- The **struct linger** parameter (defined in `<sys/socket.h>`) specifies the desired state of the option in the `l_onoff` field and the linger interval in the `l_linger` field, in seconds. A value of 0 causes a reset on the socket when the application closes the socket.
- SO\_OOBINLINE** *level: SOL\_SOCKET*
- For protocols that support out-of-band data, allows or disallows out-of-band data to be placed in the normal data input queue as received. The data is accessible using the *recv()* or *read()* calls without the MSG\_OOB flag. Some protocols always behave as if this option is set.
- SO\_RCVBUF  
and  
SO\_SNDBUF** *level: SOL\_SOCKET*
- Gets or sets the normal buffer sizes allocated for output (SO\_SNDBUF) and input (SO\_RCVBUF) buffers. You can increase the buffer size for high-volume connections, or decrease it to limit the possible backlog

of incoming data. The system places an absolute limit on these values and defaults them to at least 16K for TCP sockets.

*level:* SOL\_SOCKET

**SO\_RCVLOWAT**

Gets or sets the minimum count for input operations (default is 1). In general, receive calls block until any (nonzero) amount of data is received, and then return with the amount available or the amount requested, whichever is smaller.

If you set the value to be larger than the default, blocking receive calls will wait until they've received the low-water mark value or the requested amount, whichever is smaller. Receive calls may still return less than the low-water mark if: an error occurs, a signal is caught, or if the type of data next in the receive queue differs from that returned.

*level:* SOL\_SOCKET

**SO\_RCVTIMEO**

Gets or sets a timeout value for input operations. It accepts a **struct timeval** parameter (defined in `<sys/time.h>`) with the number of seconds and microseconds used to limit waits for input operations to complete.

In the current implementation, this timer is restarted each time additional data is received by the protocol, so the limit is in effect an inactivity timer. If a receive operation has been blocked for this much time without receiving additional data, it returns with a short count or, if no data was received, with the error EWOULDBLOCK.

*level:* SOL\_SOCKET

**SO\_REUSEADDR**

Enables or disables the reuse of duplicate addresses and port bindings. Indicates that the rules used in validating addresses supplied in a *bind()* call allows/disallows local addresses to be reused.

*level:* SOL\_SOCKET

**SO\_REUSEPORT**

Enables or disables duplicate address and port bindings. Complete duplicate bindings by multiple processes are allowed when they all set SO\_REUSEPORT before binding the port. This option permits multiple instances of a program to each receive UDP/IP multicast or broadcast



datagrams destined for the bound port. See the `reuseport_unicast` option of the `npm-tcpip.so` utility to see how unicast packets are also received on all sockets bound to the same port.

*level:* SOL\_SOCKET

**SO\_SNDLOWAT**

Gets or sets the minimum count for output operations. In BSD, this count is typically 2048, but it is a calculated value in Neutrino. If you require a specific `SO_SNDLOWAT`, you must specify the count. Most output operations process all of the data supplied by the call, delivering data to the protocol for transmission and blocking as necessary for flow control. Nonblocking output operations will process as much data as permitted (subject to flow control without blocking), but will process no data if flow control doesn't allow the smaller of the low-water mark value or the entire request to be processed.

A `select()` operation that tests the ability to write to a socket returns true only if the low-water mark amount could be processed.

*level:* SOL\_SOCKET

**SO\_SNDTIMEO**

Gets or sets a timeout value for output operations. It accepts a `struct timeval` parameter (defined in `<sys/time.h>`) that includes the number of seconds and microseconds that are used to limit waits for output operations to complete. If a send operation has blocked for this much time, it returns with a partial count or with the error `EWOULDBLOCK` if data weren't sent.

This timer is restarted each time additional data is delivered to the protocol, implying that the limit applies to output portions ranging in size from the low-water mark to the high-water mark for output. Timeouts are restricted to 32 seconds or under.

*level:* SOL\_SOCKET

**SO\_TIMESTAMP**

Enables or disables the reception of a timestamp with datagrams. If enabled on a `SOCK_DGRAM` socket, the `recvmsg()` call returns a timestamp corresponding to when the datagram was received. The `msg_control` field in the `msg_hdr` structure points to a buffer that

contains a `cmsghdr` structure followed by a `struct timeval`. The `cmsghdr` fields have the following values:

```
cmsg_len = sizeof(struct cmsghdr) + sizeof(struct timeval)
cmsg_level = SOL_SOCKET
cmsg_type = SCM_TIMESTAMP
```

**SO\_TYPE** *level:* SOL\_SOCKET

Applies to `getsockopt()` only.

Gets the type of the socket (e.g. `SOCK_STREAM`). This information is useful for servers that inherit sockets on startup.

*level:* SOL\_SOCKET

**SO\_USELOOPBACK**

Enables or disables the sending process to receive its own routing messages.

*level:* IPPROTO\_TCP

**TCP\_KEEPALIVE**

Gets or sets the amount of time in seconds between keepalive probes (the default value is 2 hours). It accepts a `struct timeval` parameter with the number of seconds to wait between the keepalive probes.

*level:* IPPROTO\_TCP

**TCP\_NODELAY**

Don't delay sending in order to coalesce packets. Under most circumstances, TCP sends data when it's presented. When outstanding data hasn't yet been acknowledged, TCP gathers small amounts of output to be sent in a single packet once an acknowledgment is received.




---

For a few clients (such as windowing systems that send a stream of mouse events that receive no replies), this packetization may cause significant delays. Therefore, TCP provides a boolean option, `TCP_NODELAY`, to defeat this algorithm.

---

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EBADF Invalid file descriptor *s*.
- EDOM Value was set out of range.
- EFAULT The address pointed to by *optval* isn't in a valid part of the process address space. For *getsockopt()*, this error may also be returned if *optlen* isn't in a valid part of the process address space.
- EINVAL The *optval* argument can't be NULL; *optlen* can't be 0.
- ENOPROTOOPT  
The option is unknown at the level indicated.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**


---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

ICMP, IP, TCP, and UDP protocols

*close()*, *getprotobyname()*, *ioctl()*, *read()*, *select()*, *setsockopt()*,  
*socket()*

*/etc/protocols* in the *Utilities Reference*

### **Synopsis:**

```
#include <sys/types.h>
#include <shadow.h>

struct spwd* getspent(void);

struct spwd* getspent_r(struct spwd* result,
 char* buffer,
 int buflen);
```

### **Arguments:**

These arguments apply only to *getspent\_r()*:

*result*     A pointer to a **spwd** structure where the function can store the entry. For more information about this structure, see *putspent()*.

*buffer*     A block of memory that the function can use to allocate storage referenced by the **spwd** structure. You can determine the maximum size needed for this buffer by calling *sysconf()* with an argument of `_SC_GETPW_R_SIZE_MAX`.

*bufsize*     The size of the block that *buffer* points to, in characters.

### **Library:**

`libc`

### **Description:**

The *getspent()* and *getspent\_r()* functions return the next entry from the shadow password database. The *getspent()* function uses a static buffer that's overwritten by each call.



---

The *fgetspent()*, *getspent()*, *getspnam()*, and functions share the same static buffer.

---

## Returns:

The *getspent()* function returns a pointer to an object of type **struct spwd** containing the next entry from the shadow password database. When *getspent()* is first called, the database is opened, and remains open until either a NULL is returned to signify end-of-file, or *endspent()* is called.

## Errors:

The *getspent()* function uses the following functions, and as a result, *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <pwd.h>
#include <shadow.h>

/*
 * This program loops, reading a login name from standard
 * input and checking to see if it is a valid name. If it
 * is not valid, the entire contents of the name in the
 * password database are printed.
 */

int main(int argc, char** argv)
{
 struct spwd* sp;
```

```

char buf[80];

setpwent ();
while(gets(buf) != NULL) {
 if((sp = getspnam(buf)) != (struct spwd *)0) {
 printf("Valid login name is: %s\n", sp->sp_namp);
 } else {
 setspent ();
 while((sp=getspent()) != (struct spwd *)0)
 printf("%s\n", sp->sp_namp);
 }
}
endspent ();
return(EXIT_SUCCESS);
}

```

**Classification:**

Unix

***getspent()*****Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

***getspent\_r()*****Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, fgetspent(), getgrent(), getlogin(), getsnam(), getpwuid(),  
putspent(), setspent()*



## Synopsis:

```
#include <sys/types.h>
#include <shadow.h>

struct spwd* getspnam(char* name);

struct spwd* getspnam_r(const char* name,
 struct spwd* result,
 char* buffer,
 size_t bufsize);
```

## Arguments:

- name*        The name of the user.
- result*      (*getspnam\_r()* only) A pointer to a **spwd** structure where the function can store the entry. For more information about this structure, see *putspent()*.
- buffer*      (*getspnam\_r()* only) A block of memory that the function can use to allocate storage referenced by the **spwd** structure. You can determine the maximum size needed for this buffer by calling *sysconf()* with an argument of `_SC_GETPW_R_SIZE_MAX`.
- bufsize*     (*getspnam\_r()* only) The size of the block that *buffer* points to, in characters.

## Library:

**libc**

## Description:

The *getspnam()* and *getspnam\_r()* functions allow a process to gain more knowledge about a user *name*. The *getspnam()* function uses a static buffer that's overwritten by each call.



The *fgetspent()*, *getspent()*, and *getspnam()* functions share the same static buffer.

## Returns:

A pointer to an object of type `struct spwd` containing an entry from the group database with a matching *name*, or NULL if an error occurred or the function couldn't find a matching entry.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <pwd.h>
#include <shadow.h>

/*
 * Print information from the password entry
 * about the user name given as argv[1].
 */

int main(int argc, char** argv)
{
 struct spwd* sp;

 if (argc < 2) {
 printf("%s username \n", argv[0]);
 return(EXIT_FAILURE);
 }

 if((sp = getspnam(argv[1])) == (struct spwd*)0) {
 fprintf(stderr, "getspnam: unknown %s\n",
 argv[1]);
 return(EXIT_FAILURE);
 }
 printf("login name %s\n", sp->sp_namp);
 printf("password %s\n", sp->sp_pwdp);
 return(EXIT_SUCCESS);
}
```

**Classification:**

*getspnam()* is POSIX 1003.1; *getspnam\_r()* is Unix

***getspnam()***

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

***getspnam\_r()***

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fgetspent()*, *getlogin()*, *getspent()*, *getpwuid()*, *putspent()*

## ***getsubopt()***

© 2004, QNX Software Systems Ltd.

*Parse suboptions from a string*

### **Synopsis:**

```
#include <stdlib.h>

int getsubopt(char** optionp,
 char* const* tokens,
 char** valuep);
```

### **Arguments:**

*optionp*     The address of a pointer to the string of options that you want to parse. The function updates this pointer as it parses the options; see below.

*tokens*     A vector of possible tokens.

*valuep*     The address of a pointer that the function updates to point to the first character of a value that's associated with an option; see below.

### **Library:**

`libc`

### **Description:**

The *getsubopt()* functions parses suboptions in a flag argument that was initially parsed by *getopt()*. These suboptions are separated by commas and may consist of either a single token or a token-value pair separated by an equal sign. Since commas delimit suboptions in the option string, they aren't allowed to be part of the suboption or the value of a suboption. A command that uses this syntax is **mount**, which allows the user to specify mount parameters with the **-o** option as follows:

```
mount -o rw,hard,bg,wsiz=1024 speed:/usr /usr
```

In this example there are four suboptions: **rw**, **hard**, **bg**, and **wsiz**, the last of which has an associated value of 1024.

The *getsubopt()* function takes the address of a pointer to the option string, a vector of possible tokens, and the address of a value string pointer. It returns the index of the token that matched the suboption in the input string, or -1 if there was no match. If the option string at *optionp* contains only one suboption, *getsubopt()* updates *optionp* to point to the null character at the end of the string; otherwise, it isolates the suboption by replacing the comma separator with a null character, and updates *optionp* to point to the start of the next suboption. If the suboption has an associated value, *getsubopt()* updates *valuep* to point to the value's first character. Otherwise, it sets *valuep* to NULL.

The token vector is organized as a series of pointers to null strings. The end of the token vector is identified by a NULL pointer.

When *getsubopt()* returns, if *valuep* isn't NULL, the suboption processed included a value. The calling program may use this information to determine if the presence or lack of a value for this suboption is an error.

Additionally, when *getsubopt()* fails to match the suboption with the tokens in the *tokens* array, the calling program should decide if this is an error, or if the unrecognized option should be passed to another program.

## Returns:

The *getsubopt()* function returns -1 when the token it's scanning isn't in the *tokens* vector. The variable addressed by *valuep* contains a pointer to the first character of the token that wasn't recognized rather than a pointer to a value for that token.

The variable addressed by *optionp* points to the next option to be parsed, or a null character if there are no more options.

## Examples:

The following code fragment shows how to process options to the `mount(1M)` command using *getsubopt()*:

```
#include <stdlib.h>
```

```
char *myopts[] = {
#define READONLY 0
 "ro",
#define READWRITE 1
 "rw",
#define WRITESIZE 2
 "wsize",
#define READSIZE 3
 "rsize",
 NULL};

main(argc, argv)
int argc;
char **argv;
{
 int sc, c, errflag;
 char *options, *value;
 extern char *optarg;
 extern int optind;
 .
 .
 .
 while((c = getopt(argc, argv, "abf:o:")) != -1) {
 switch (c) {
 case 'a': /* process a option */
 break;
 case 'b': /* process b option */
 break;
 case 'f':
 ofile = optarg;
 break;
 case '?':
 errflag++;
 break;
 case 'o':
 options = optarg;
 while (*options != '\0') {
 switch(getsubopt(&options, myopts, &value)) {
 case READONLY : /* process ro option */
 break;
 case READWRITE : /* process rw option */
 break;

 case WRITESIZE : /* process wsize option */
 if (value == NULL) {
 error_no_arg();
 errflag++;
 } else
 write_size = atoi(value);
 break;
 }
 }
 }
 }
}
```

```
 case READSIZE : /* process rsize option */
 if (value == NULL) {
 error_no_arg();
 errflag++;
 } else
 read_size = atoi(value);
 break;
 default :
 /* process unknown token */
 error_bad_token(value);
 errflag++;
 break;
 }
}
break;
}
}
if (errflag) {
 /* print usage instructions etc. */
}
for (; optind < argc; optind++) {
 /* process remaining arguments */
}
.
.
.
}
```

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

During parsing, commas in the option input string are changed to null characters. White space in tokens or token-value pairs must be protected from the shell by quotes.

**See also:**

*getopt()*



## Synopsis:

```
#include <sys/time.h>

int gettimeofday(struct timeval * when,
 void * not_used);
```

## Arguments:

- when*            A pointer to a `timeval` structure where the function can store the time. The `struct timeval` contains the following members:
- `long tv_sec` — the number of seconds since the start of the Unix Epoch.
  - `long tv_usec` — the number of microseconds.
- not\_used*        This pointer must be NULL or the behavior of `gettimeofday()` is unspecified. This argument is provided only for backwards compatibility.

## Library:

`libc`

## Description:

The `gettimeofday()` function returns the current time in *when* in seconds and microseconds, since the Unix Epoch, 00:00:00 January 1, 1970 Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time (GMT)).

## Returns:

0 for success, or -1 if an error occurs (*errno* is set).

**Errors:**

EFAULT      An error occurred while accessing the *when* buffer.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The *gettimeofday()* function is provided for porting existing code. You shouldn't use it in new code; use *clock\_gettime()* instead.

**See also:**

*asctime()*, *asctime\_r()*, *clock\_gettime()*, *clock\_settime()*, *ctime()*, *ctime\_r()*, *difftime()*, *gmtime()*, *gmtime\_r()*, *localtime()*, *localtime\_r()*, *settimeofday()*, *time()*

## Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

uid_t getuid(void);
```

## Library:

libc

## Description:

The *getuid()* function gets the user ID for the calling process.

## Returns:

The user ID for the calling process.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main(void)
{
 printf("My userid is %d\n", getuid());
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*getegid(), geteuid(), getgid(), setuid()*

**Synopsis:**

```
#include <utmp.h>

struct utmp * getutent(void);
```

**Library:**

libc

**Description:**

The *getutent()* function reads in the next entry from a user-information file. If the file isn't already open, *getutent()* opens it. If the function reaches the end of the file, it fails.

**Returns:**

A pointer to a **utmp** structure for the next entry, or NULL if the file couldn't be read or reached the end of file.

**Files:**

*PATH\_UTMP*  
Specifies the user information file.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## **Caveats:**

The most current entry is saved in a static structure. Copy it before making further accesses.

On each call to either *getutid()* or *getutline()*, the routine examines the static structure before performing more I/O. If the contents of the static structure match what it's searching for, the function looks no further. For this reason, to use *getutline()* to search for multiple occurrences, zero out the static area after each success, or *getutline()* will return the same structure over and over again.

There's one exception to the rule about emptying the structure before further reads are done: the implicit read done by *pututline()* (if it finds that it isn't already at the correct place in the file) doesn't hurt the contents of the static structure returned by the *getutent()*, *getutid()* or *getutline()* routines, if you just modified those contents and passed the pointer back to *pututline()*.

These routines use buffered standard I/O for input, but *pututline()* uses an unbuffered nonstandard write to avoid race conditions between processes trying to modify the **utmp** and **wtmp** files.

## **See also:**

*endutent()*, *getutid()*, *getutline()*, *pututline()*, *setutent()*, **utmp**, *utmpname()*

**login** in the *Utilities Reference*

## Synopsis:

```
#include <utmp.h>

struct utmp * getutid(struct utmp * id);
```

## Arguments:

*id* A pointer to a **utmp** structure that you want to find in the user-information file.

## Library:

**libc**

## Description:

The *getutid()* function searches forward from the current point in the **utmp** file until it finds a matching entry:

- If *id->ut\_type* is one of **RUN\_LVL**, **BOOT\_TIME**, **OLD\_TIME**, or **NEW\_TIME**, the function looks for an entry with the same *ut\_type*.
- If *id->ut\_type* is **INIT\_PROCESS**, **LOGIN\_PROCESS**, **USER\_PROCESS**, or **DEAD\_PROCESS**, *getutid()* looks for the first entry entry with the same *ut\_type* and a *ut\_id* field that matches *id->ut\_id*.

If *getutid()* reaches the end of the file without finding a match, the search fails.

## Returns:

A pointer to the **utmp** structure for the matching entry, or **NULL** if it couldn't be found.

**Files:**

*\_PATH\_UTMP*

Specifies the user information file.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

The most current entry is saved in a static structure. Copy it before making further accesses.

On each call to either *getutid()* or *getutline()*, the routine examines the static structure before performing more I/O. If the contents of the static structure match what it's searching for, the function looks no further. For this reason, to use *getutline()* to search for multiple occurrences, zero out the static area after each success, or *getutline()* will return the same structure over and over again.

There's one exception to the rule about emptying the structure before further reads are done: the implicit read done by *pututline()* (if it finds that it isn't already at the correct place in the file) doesn't hurt the contents of the static structure returned by the *getutent()*, *getutid()* or *getutline()* routines, if the user has just modified those contents and passed the pointer back to *pututline()*.



These routines use buffered standard I/O for input, but *pututline()* uses an unbuffered nonstandard write to avoid race conditions between processes trying to modify the **utmp** and **wtmp** files.

**See also:**

*endutent()*, *getutent()*, *getutline()*, *pututline()*, *setutent()*, **utmp**, *utmpname()*

**login** in the *Utilities Reference*

## ***getutline()***

© 2004, QNX Software Systems Ltd.

*Get an entry from the user-information file*

---

### **Synopsis:**

```
#include <utmp.h>

struct utmp * getutline(struct utmp * line);
```

### **Arguments:**

*line*     A pointer to a **utmp** structure that you want to find in the user-information file.

### **Library:**

**libc**

### **Description:**

The *getutline()* function searches forward from the current point in the **utmp** file until it finds an entry of the type **LOGIN\_PROCESS** or a *ut\_line* string that matches *line->ut\_line*. If the function reaches the end of the file is reached without finding a match, the function fails.

### **Returns:**

A pointer to the **utmp** structure for the entry found, or **NULL** if the search failed.

### **Files:**

*\_PATH\_UTMP*  
Specifies the user information file.

### **Classification:**

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### Caveats:

The most current entry is saved in a static structure. Copy it before making further accesses.

On each call to either *getutid()* or *getutline()*, the routine examines the static structure before performing more I/O. If the contents of the static structure match what it's searching for, the function looks no further. For this reason, to use *getutline()* to search for multiple occurrences, zero out the static area after each success, or *getutline()* will return the same structure over and over again.

There's one exception to the rule about emptying the structure before further reads are done: the implicit read done by *pututline()* (if it finds that it isn't already at the correct place in the file) doesn't hurt the contents of the static structure returned by the *getutent()*, *getutid()* or *getutline()* routines, if the user has just modified those contents and passed the pointer back to *pututline()*.

These routines use buffered standard I/O for input, but *pututline()* uses an unbuffered nonstandard write to avoid race conditions between processes trying to modify the **utmp** and **wtmp** files.

### See also:

*endutent()*, *getutent()*, *getutid()*, *pututline()*, *setutent()*, **utmp**, *utmpname()*

**login** in the *Utilities Reference*

## ***getw()***

© 2004, QNX Software Systems Ltd.

*Get a word from a stream*

---

### **Synopsis:**

```
#include <stdio.h>

int getw(FILE* stream);
```

### **Arguments:**

*stream*     The stream that you want to read a word from.

### **Library:**

`libc`

### **Description:**

The *getw()* function returns the next word (i.e. integer) from the named input stream. This function increments the associated file pointer, if defined, to point to the next word. The size of a word is the size of an integer, and varies from machine to machine. The *getw()* function assumes no special alignment in the file.

### **Returns:**

The next word, or the constant EOF at the end-of-file or on an error; it sets the EOF or error indicator of the stream.



Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

---

### **Errors:**

`E_OVERFLOW`     The file is a regular file, and an attempt was made to read at or beyond the offset maximum associated with the corresponding stream.

**Classification:**

Legacy Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Because of possible differences in word length and byte ordering, files written using *putw()* are implementation-dependent, and might not be read correctly using *getw()* on a different processor.

**See also:**

*fclose()*, *feof()*, *ferror()*, *fgetc()*, *flockfile()*, *fopen()*, *fread()*, *getc()*, *getc\_unlocked()*, *getchar()*, *getchar\_unlocked()*, *gets()*, *putc()*, *putw()*, *scanf()*, *ungetc()*,

## ***getwc()***

© 2004, QNX Software Systems Ltd.

*Read a wide character from a stream*

### **Synopsis:**

```
#include <wchar.h>

wint_t getwc(FILE * fp);
```

### **Arguments:**

*fp* The stream from which you want to read a wide character.

### **Library:**

libc

### **Description:**

The *getwc()* function reads the next wide character from the specified stream.

### **Returns:**

The next character from the stream, cast as (**wint\_t**) (**wchar\_t**), or WEOF if end-of-file has been reached or if an error occurs (*errno* is set).



---

Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

---

### **Errors:**

|        |                                                                                         |
|--------|-----------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for <i>fp</i> and would have been blocked by this operation. |
| EBADF  | The <i>fp</i> stream isn't valid for reading.                                           |
| EINTR  | A signal terminated the read operation; no data was transferred.                        |

|            |                                                                                                                |
|------------|----------------------------------------------------------------------------------------------------------------|
| EIO        | Either a physical I/O error has occurred, or the process is in the background and is being ignored or blocked. |
| E_OVERFLOW | Cannot read at or beyond the offset maximum for this stream.                                                   |

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*errno*, *feof()*, *ferror()*, *putwc()*, *putwchar()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter

## ***getwchar()***

© 2004, QNX Software Systems Ltd.

*Read a character from a stream*

### **Synopsis:**

```
#include <wchar.h>

wint_t getwchar(void);
```

### **Library:**

libc

### **Description:**

The *getwchar()* function reads the next wide character from *stdin*.

### **Returns:**

The next character from *stdin*, cast as (**wint\_t**) (**wchar\_t**), or WEOF if the end-of-file has been reached or if an error occurs (*errno* is set).



---

Use *feof()* or *ferror()* to distinguish an end-of-file condition from an error.

---

### **Errors:**

|            |                                                                                                                |
|------------|----------------------------------------------------------------------------------------------------------------|
| EAGAIN     | The O_NONBLOCK flag is set for <i>stdin</i> and would have been blocked by this operation.                     |
| EINTR      | A signal terminated the read operation; no data was transferred.                                               |
| EIO        | Either a physical I/O error has occurred, or the process is in the background and is being ignored or blocked. |
| E_OVERFLOW | Cannot read at or beyond the offset maximum for this stream.                                                   |



## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *feof()*, *ferror()*, *putwc()*, *putwchar()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter

## ***getwd()***

© 2004, QNX Software Systems Ltd.

*Get current working directory pathname*

### **Synopsis:**

```
#include <unistd.h>

char* getwd(char *path_name);
```

### **Arguments:**

*path\_name*     A buffer where the function can store the current working directory.

### **Library:**

libc

### **Description:**

The *getwd()* function determines the absolute pathname of the current working directory of the calling process, and copies that pathname into the array pointed to by the *path\_name* argument.

If the length of the pathname of the current working directory is greater than (`{PATH_MAX} + 1`) including the null byte, *getwd()* fails and returns a null pointer.



---

For portability, use *getcwd()* instead of *getwd()*.

---

### **Returns:**

A pointer to the string containing the absolute pathname of the current working directory. On error, *getwd()* returns a null pointer and the contents of the array pointed to by *path\_name* are undefined.

### **Classification:**

Legacy Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*getcwd()*

# **glob()**

© 2004, QNX Software Systems Ltd.

Find paths matching a pattern

## **Synopsis:**

```
#include <glob.h>

int glob(const char* pattern,
 int flags,
 int (*errfunc)(const char* epath,
 int error),
 glob_t* pglob);
```

## **Arguments:**

- pattern* The pattern you want to match. This can include these wildcard characters:
- \* matches any string, of any length
  - ? matches any single character
  - [chars] matches any of the characters found in the string *chars*.
- flags* Flags that affect the search; see below.
- errfunc* A pointer to a function that *glob()* calls when it encounters a directory that it can't open or read. For more information, see below.
- pglob* A pointer to a **glob\_t** structure where *glob()* can store the paths found. This structure contains at least the following members:
- **size\_t** *gl\_pathc* — the number of pathnames matched by *pattern*.
  - **char\*\*** *gl\_pathv* — a NULL-terminated array of pointers to the pathnames matched by *pattern*.
  - **size\_t** *gl\_offs* — the number of pointers to reserve at the beginning of *gl\_pathv*.

You must create the **glob\_t** structure before calling *glob()*. The *glob()* function allocates storage as needed for the *gl\_pathv* array. Use *globfree()* to free this space.

## Library:

`libc`

## Description:

The *glob()* function finds pathnames matching the given pattern.

In order to have access to a pathname, *glob()* must have search permission on every component of the path except the last, and read permission on each directory of every filename component of *pattern* that contains any of the special characters (`*`, `?`, `[` and `]`).

The *errfunc* argument is a pointer to an error-handler function with this prototype:

```
int errfunc(const char* epath, int error);
```

The *errfunc* function is called when *glob()* encounters a directory that it can't open or read. The arguments are:

*epath*     A pointer to the path that failed.

*error*     The value of *errno* from the failure. The *error* argument can be set to any of the values returned by *opendir()*, *readdir()*, or *stat()*.

The *errfunc* function should return 0 if *glob()* should continue, or a nonzero value if *glob()* should stop searching.

You can set *errfunc* to NULL to ignore these types of errors.

The *flags* argument can be set to any combination of the following bits:

GLOB\_APPEND     Append found pathnames to the ones from a previous call from *glob()*.

GLOB\_DOOFFS     Use the value in *pglob->gl\_offs* to specify how many NULL pointers to add at the beginning of *pglob->pathv*. After the call to *glob()*,

*pglob->pathv* will contain *pglob->gl\_offs* NULL pointers, followed by *pglob->gl\_pathc* pathnames, followed by a NULL pointer. This can be useful if you're building a command to be applied to the matched files.

- GLOB\_ERR** Cause *glob()* to return when it encounters a directory that it can't open or read. Otherwise, *glob()* will continue to find matches.
- GLOB\_MARK** Append a slash to each matching pathname that's a directory.
- GLOB\_NOCHECK**  
If *pattern* doesn't match any path names, return only the contents of *pattern*.
- GLOB\_NOESCAPE**  
Disable backslash escapes in *pattern*.
- GLOB\_NOSORT** Don't sort the returned pathnames; their order will be unspecified. The default is to sort the pathnames.

## Returns:

Zero for success, or an error value.

## Errors:

- GLOB\_ABEND**  
The scan was stopped because **GLOB\_ERR** was set, or the *errfunc* function returned nonzero.
- GLOB\_NOMATCH**  
The value of *pattern* doesn't match any existing pathname, and **GLOB\_NOCHECK** wasn't set in *flags*.
- GLOB\_NOSPACE**  
Unable to allocate memory to store the matched paths.

**Examples:**

This simple example attempts to find all of the `.c` files in the current directory and print them in the order the filesystem found them.

```
#include <unistd.h>
#include <stdio.h>
#include <glob.h>

int main(void)
{
 glob_t paths;
 int retval;

 paths.gl_pathc = 0;
 paths.gl_pathv = NULL;
 paths.gl_offs = 0;

 retval = glob("*.c", GLOB_NOCHECK | GLOB_NOSORT,
 NULL, &paths);
 if(retval == 0) {
 int idx;

 for(idx = 0; idx < paths.gl_pathc; idx++) {
 printf("[%d]: %s\n", idx,
 paths.gl_pathv[idx]);
 }

 globfree(&paths);
 } else {
 puts("glob() failed");
 }

 return 0;
}
```

**Classification:**

POSIX 1003.1a

**Safety**

Cancellation point    Yes

Interrupt handler    No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

Don't change the values in *pglob* between calling *glob()* and *globfree()*.

**See also:**

*globfree()*, *wordexp()*, *wordfree()*



**Synopsis:**

```
#include <glob.h>

void globfree(glob_t* pglob);
```

**Arguments:**

*pglob* A pointer to a `glob_t` structure that you passed to `glob()`.

**Library:**

`libc`

**Description:**

The `globfree()` function frees the storage allocated by a call to `glob()`.

**Classification:**

POSIX 1003.1a

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Don't change the values in *pglob* between calling `glob()` and `globfree()`.

**See also:**

*glob()*, *wordexp()*, *wordfree()*

## Synopsis:

```
#include <time.h>

struct tm* gmtime(const time_t* timer);
```

## Arguments:

*timer*     A pointer to a `time_t` structure that contains the time that you want to convert.

## Library:

`libc`

## Description:

The *gmtime()* function converts the calendar time pointed to by *timer* into a broken-down time, expressed as Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time or GMT).

The *gmtime()* function places the converted time in a static structure that's reused each time you call *gmtime()*. If you want a thread-safe version of *gmtime()*, use *gmtime\_r()*.

You typically use the `date` command to set the computer's internal clock using Coordinated Universal Time (UTC). Use the **TZ** environment variable or `_CS_TIMEZONE` configuration string to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

## Returns:

A pointer to a `tm` structure that contains the broken-down time.

**Classification:**

ANSI, POSIX 1003.1

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*asctime(), asctime\_r(), clock(), ctime(), difftime(), gmtime\_r(), localtime(), localtime\_r(), mktime(), strftime(), time(), tm, tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*

**Synopsis:**

```
#include <time.h>

struct tm* gmtime_r(const time_t* timer,
 struct tm* result);
```

**Arguments:**

*timer*     A pointer to a `time_t` structure that contains the time that you want to convert.

*result*    A pointer to a `tm` structure where the function can store the broken-down time.

**Library:**

`libc`

**Description:**

The `gmtime_r()` function converts the calendar time pointed to by *timer* into a broken-down time, expressed as Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time or GMT) and stores it in the `tm` structure pointed to by *result*.

You typically use the `date` command to set the computer's internal clock using Coordinated Universal Time (UTC). Use the `TZ` environment variable or `_CS_TIMEZONE` configuration string to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

**Returns:**

A pointer to the `tm` structure containing the broken-down time.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*asctime(), asctime\_r(), clock(), ctime(), difftime(), localtime(), localtime\_r(), mktime(), strftime(), time(), tm, tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*

## Synopsis:

```
#include <netdb.h>

extern int h_errno;
```

## Library:

libsocket

## Description:

The *h\_errno* variable can be set by any one of the following functions:

- *gethostbyaddr()*
- *gethostbyaddr\_r()*
- *gethostbyname()*
- *gethostbyname2()*
- *gethostbyname\_r()*
- *res\_query()*
- *res\_search()*

It can be set to any one of the following:

HOST\_NOT\_FOUND

Authoritative answer: Unknown host.

NETDB\_INTERNAL

You specified an invalid address family when calling *gethostbyname2()*.

NO\_DATA

Valid name, no data record of the requested type. The name is known to the name server, but has no IP address associated with it — this isn't a temporary error. Another type of request to the

name server using this domain name will result in an answer (e.g. a mail-forwarder may be registered for this domain).

**NO\_RECOVERY**

Unknown server error. An unexpected server failure was encountered. This is a nonrecoverable network error.

**TRY\_AGAIN**

Nonauthoritative answer: Host name lookup failure. This is usually a temporary error and means that the local server didn't receive a response from an authoritative server. A retry at some later time may succeed.

**Classification:**

POSIX 1003.1-2001

**Caveats:**

Unlike *errno*, *h\_errno* isn't thread-safe.

**See also:**

*errno*, *gethostbyaddr()*, *gethostbyaddr\_r()*, *gethostbyname()*, *gethostbyname2()*, *gethostbyname\_r()*, *res\_query()*, *res\_search()*



## Synopsis:

```
#include <search.h>

int hcreate(size_t nel);
```

## Arguments:

*nel* An estimate of the maximum number of entries that the table will contain. The algorithm might adjust this number upward in order to obtain certain mathematically favorable circumstances.

## Library:

`libc`

## Description:

The *hcreate()* function allocates space for the hash search table. You must call this function before using *hsearch()*.

The *hsearch()* and *hcreate()* functions use *malloc()* to allocate space.

Only one hash search table may be active at any given time. You can destroy the table by calling *hdestroy()*.

## Returns:

0 if there isn't enough space available to allocate the table.

## Examples:

See *hsearch()*.

## Classification:

Standard Unix

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## **See also:**

*bsearch()*, *hdestroy()*, *hsearch()*, *malloc()*

*The Art of Computer Programming, Volume 3, Sorting and Searching*  
by Donald E. Knuth, published by Addison-Wesley Publishing  
Company, 1973.

**Synopsis:**

```
#include <search.h>

void hdestroy(void);
```

**Library:**

libc

**Description:**

The *hdestroy()* function destroys the hash search table that was created by *hcreate()* and used by *hsearch()*. Only one hash search table may be active at any given time.

**Examples:**

See *hsearch()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*hcreate()*, *hsearch()*

## ***herror()***

© 2004, QNX Software Systems Ltd.

*Print the message associated with the value of `h_errno` to standard error*

### **Synopsis:**

```
#include <netdb.h>

void herror(const char* prefix);
```

### **Arguments:**

*prefix*    NULL, or a string that you want to print before the error message.

### **Library:**

`libsocket`

### **Description:**

The *herror()* function prints the message corresponding to the error number contained in *h\_errno* to *stderr*. The following functions can set *h\_errno*:

- *gethostbyaddr()*
- *gethostbyaddr\_r()*
- *gethostbyname()*
- *gethostbyname2()*
- *gethostbyname\_r()*
- *res\_query()*
- *res\_search()*

If the *prefix* string is non-NULL, it's printed, followed by a colon and a space. The error message is printed with a trailing newline. One of the following messages could be printed:

HOST\_NOT\_FOUND

Authoritative answer: Unknown host.

**NETDB\_INTERNAL**

You specified an invalid address family when calling *gethostbyname2()*.

**NO\_DATA**

Valid name, no data record of the requested type. The name is known to the name server, but has no IP address associated with it — this isn't a temporary error. Another type of request to the name server using this domain name will result in an answer (e.g. a mail-forwarder may be registered for this domain).

**NO\_RECOVERY**

Unknown server error. An unexpected server failure was encountered. This is a nonrecoverable network error.

**TRY\_AGAIN**

Nonauthoritative answer: Host name lookup failure. This is usually a temporary error and means that the local server didn't receive a response from an authoritative server. A retry at some later time may succeed.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*gethostbyaddr(), gethostbyaddr\_r(), gethostbyname(),  
gethostbyname\_r(), h\_errno, res\_query(), res\_search(), stderr*

**Synopsis:**

```
#include <netdb.h>

struct hostent {
 char * h_name;
 char ** h_aliases;
 int h_addrtype;
 int h_length;
 char ** h_addr_list;
};

#define h_addr h_addr_list[0]
```

**Description:**

This structure describes an Internet host. It contains either the information obtained from a name server, or broken-out fields from a line in `/etc/hosts`.

The members of this structure are:

|                    |                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------|
| <i>h_name</i>      | The official name of the host.                                                                                |
| <i>h_aliases</i>   | A zero-terminated array of alternate names for the host.                                                      |
| <i>h_addrtype</i>  | The type of address being returned; currently always AF_INET.                                                 |
| <i>h_length</i>    | The length of the address, in bytes.                                                                          |
| <i>h_addr_list</i> | A zero-terminated array of network addresses for the host. Host addresses are returned in network byte order. |

A `#define` statement is used to define the following:

|               |                                                                               |
|---------------|-------------------------------------------------------------------------------|
| <i>h_addr</i> | The first address in <i>h_addr_list</i> . This is for backward compatibility. |
|---------------|-------------------------------------------------------------------------------|

**Classification:**

Unix, POSIX 1003.1-2001

**See also:**

*endhostent()*, *gethostbyaddr()*, *gethostbyname()*, *gethostent()*,  
*sethostent()*

*/etc/hosts*, */etc/resolv.conf* in the *Utilities Reference*



**Synopsis:**

```
#include <search.h>

ENTRY* hsearch (ENTRY item,
 ACTION action);
```

**Arguments:**

- item* A structure of type **ENTRY**, defined in `<search.h>`, that contains:
- **char \*key** — a pointer to the comparison key.
  - **void \*data** — a pointer to any other data to be associated with the key.
- action* A member of an enumeration type **ACTION**, also defined in `<search.h>`, indicating what to do with the entry if it isn't in the table:
- **ENTER** — insert the entry in the table at the appropriate point. If the item is a duplicate of an existing item, the new item isn't added, and *hsearch()* returns a pointer to the existing one.
  - **FIND** — don't add the entry. If the item can't be found, *hsearch()* returns **NULL**.

**Library:**

`libc`

**Description:**

The *hsearch()* function is a hash-table search routine generalized from Knuth (6.4) Algorithm D. Before using this function, you must call *hcreate()* to create the hash table.

The *hsearch()* function returns a pointer into a hash table indicating the location at which an entry can be found. This function uses *strcmp()* as the comparison function.

The *hsearch()* and *hcreate()* functions use *malloc()* to allocate space. Only one hash search table may be active at any given time. You can destroy the table by calling *hdestroy()*.

## Returns:

A pointer to the item found, or NULL if either the action is FIND and the item wasn't found, or the action is ENTER and the table is full.

## Examples:

The following example reads in strings followed by two numbers and stores them in a hash table, discarding duplicates. It then reads in strings, finds the matching entry in the hash table and prints it.

```
#include <stdio.h>
#include <search.h>
#include <string.h>
#include <stdlib.h>
struct info { /* this is the info stored in table */
 int age, room; /* other than the key */
};
#define NUM_EMPL 5000 /* # of elements in search table */
main()
{
 /* space to store strings */
 char string_space[NUM_EMPL*20];
 /* space to store employee info */
 struct info info_space[NUM_EMPL];
 /* next avail space in string_space */
 char *str_ptr = string_space;
 /* next avail space in info_space */
 struct info *info_ptr = info_space;
 ENTRY item, *found_item;
 /* name to look for in table */
 char name_to_find[30];
 int i = 0;

 /* create table */
 (void) hcreate(NUM_EMPL);
 while (scanf("%s%d%d", str_ptr, &info_ptr->age,
 &info_ptr->room) != EOF && i++ < NUM_EMPL) {
 /* put info in structure, and structure in item */
 item.key = str_ptr;
 item.data = (void *)info_ptr;
 str_ptr += strlen(str_ptr) + 1;
 }
}
```

```

 info_ptr++;
 /* put item into table */
 (void) hsearch(item, ENTER);
 }

 /* access table */
 item.key = name_to_find;
 while (scanf("%s", item.key) != EOF) {
 if ((found_item = hsearch(item, FIND)) != NULL) {
 /* if item is in the table */
 (void)printf("found %s, age = %d, room = %d\n",
 found_item->key,
 ((struct info *)found_item->data)->age,
 ((struct info *)found_item->data)->room);
 } else {
 (void)printf("no such employee %s\n",
 name_to_find);
 }
 }
 hdestroy();
 return 0;
}

```

**Classification:**

Standard Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***bsearch(), hcreate(), hdestroy(), malloc(), strcmp()*

*The Art of Computer Programming, Volume 3, Sorting and Searching*  
by Donald E. Knuth, published by Addison-Wesley Publishing  
Company, 1973.

*Get an error message string associated with the error return status*

**Synopsis:**

```
#include <netdb.h>

const char* hstrerror(int err);
```

**Arguments:**

*err* The error code that you want to get the message for. For more information, see *h\_errno*.

**Library:**

`libsocket`

**Description:**

The *hstrerror()* function gets an error message string associated with the error return status from network host-related functions.

Network host-related functions such as the following can return the error status:

- *gethostbyaddr()*, *gethostbyaddr\_r()*
- *gethostbyname()*, *gethostbyname\_r()*
- *res\_query()*
- *res\_search()*

You can check the external integer *h\_errno* to see whether this is a temporary failure or an invalid or unknown host.

**Returns:**

A pointer to the message string affiliated with an error number.



---

Don't modify the message string.

---

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*h\_errno, herror*

*Convert a 32-bit value from host-byte order to network-byte order*

**Synopsis:**

```
#include <arpa/inet.h>

uint32_t htonl(uint32_t hostlong);
```

**Arguments:**

*hostlong*     The value that you want to convert.

**Library:**

`libc`

**Description:**

The *htonl()* function converts a 32-bit value from host-byte order to network-byte order.

You typically use this routine in conjunction with the internet addresses and ports that *gethostbyname()* and *getservent()* return.

**Returns:**

The value in network-byte order.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*gethostbyname(), getservent(), htons(), ntohl(), ntohs()*



**Synopsis:**

```
#include <arpa/inet.h>

uint16_t htons(uint16_t hostshort);
```

**Arguments:**

*hostshort*     The value that you want to convert.

**Library:**

`libc`

**Description:**

The *htons()* function converts a 16-bit value from host-byte order to network-byte order.

You typically use this routine in conjunction with the internet addresses and ports that *gethostbyname()* and *getservent()* return.

**Returns:**

The value in network-byte order.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*gethostbyname(), getservent(), htonl(), ntohl(), ntohs()*

### **Synopsis:**

```
#include <hw/sysinfo.h>

unsigned hwi_find_item(unsigned start,
 ...);
```

### **Arguments:**

- start*        Where to start the search for the given item.
- For the initial call, set this argument to HWI\_NULL\_OFF. If the item found isn't the one that you want, pass the return value from the first call to *hwi\_find\_item()* as the *start* parameter of the next call. This makes the search pick up where it left off. You can repeat this process as many times as required (the return value from the second call going into the *start* parameter of the third, etc).
- char \**        A sequence of names for identifying the item being searched.
- Terminate the sequence with a NULL pointer. The last string before the NULL is the bottom-level item name that you're looking for, the string in front of that is the name of the item that owns the bottom-level item, etc.

### **Library:**

`libc`

### **Description:**

The *hwi\_find\_item()* function finds an item in the *hwinfo* structure of the system page.

## Returns:

The offset of the item requested, or HWI\_NULL\_OFF if the item wasn't found.

## Examples:

Find the first occurrence of an item called "foobar":

```
item_off = hwi_find_item(HWI_NULL_OFF, "foobar", NULL);
```

Find the first occurrence of an item called "foobar" that's owned by "sam":

```
item_off = hwi_find_item(HWI_NULL_OFF, "sam", "foobar", NULL);
```

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*hwi\_find\_tag()*, *hwi\_off2tag()*, *hwi\_tag2off()*

"Structure of the system page" in the Customizing Image Startup Programs chapter of *Building Embedded Systems*

### **Synopsis:**

```
#include <hw/sysinfo.h>

unsigned hwi_find_tag(unsigned start,
 int curr_item,
 const char * tagname);
```

### **Arguments:**

- |                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>start</i>     | Where to start to search for the given item.<br><br>For the initial call, set this argument to HWI_NULL_OFF. If the item found isn't the one that you want, pass the return value from the first call to <i>hwi_find_tag()</i> as the <i>start</i> parameter of the next call. This makes the search pick up where it left off. You can repeat this process as many times as required (the return value from the second call going into the <i>start</i> parameter of the third, etc). |
| <i>curr_item</i> | If this argument is nonzero, the search stops at the end of the current item (i.e. the one that <i>start</i> points to). If <i>curr_item</i> is zero, the search continues until the end of the section.                                                                                                                                                                                                                                                                               |
| <i>tagname</i>   | The name of tag to search for.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

### **Library:**

`libc`

### **Description:**

The *hwi\_find\_tag()* function finds the tag named *tagname*.

## Returns:

The offset of the tag, or HWI\_NULL\_OFF if the tag wasn't found.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*hwi\_find\_item()*, *hwi\_off2tag()*, *hwi\_tag2off()*

“Structure of the system page” in the Customizing Image Startup Programs chapter of *Building Embedded Systems*

*Return a pointer to the start of a tag in the hwinfo area of the system page*

**Synopsis:**

```
#include <hw/sysinfo.h>

void * hwi_off2tag(unsigned offsect);
```

**Arguments:**

*offsect*     The offset, in bytes from the start of the *hwinfo* section, of a tag.

**Library:**

`libc`

**Description:**

The *hwi\_off2tag()* function returns a pointer to the start of the tag, given an offset.

**Returns:**

A pointer to the start of the tag.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*hwi\_find\_item()*, *hwi\_find\_tag()*, *hwi\_tag2off()*

“Structure of the system page” in the Customizing Image Startup Programs chapter of *Building Embedded Systems*



*Return the offset from the start of the hwinfo area of the system page*

**Synopsis:**

```
#include <hw/sysinfo.h>

unsigned hwi_tag2off(void *tag);
```

**Arguments:**

*tag*     A pointer to a tag in the *hwinfo* area of the system page.

**Library:**

`libc`

**Description:**

Given a pointer to the start of a tag, the *hwi\_tag2off()* function returns the offset, in bytes, from the beginning of the start of the *hwinfo* section.

**Returns:**

The offset of the tag, in bytes.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*hwi\_find\_item()*, *hwi\_find\_tag()*, *hwi\_off2tag()*

“Structure of the system page” in the Customizing Image Startup Programs chapter of *Building Embedded Systems*

*Calculate the length of the hypotenuse for a right-angled triangle*

### **Synopsis:**

```
#include <math.h>

double hypot(double x,
 double y);

float hypotf(float x,
 float y);
```

### **Arguments:**

*x, y*     The lengths of the sides that are adjacent to the right angle.

### **Library:**

`libm`

### **Description:**

These functions compute the length of the hypotenuse for a right triangle whose sides are *x* and *y* adjacent to the right angle. The calculation is equivalent to:

```
length = sqrt(x*x + y*y);
```

### **Returns:**

The length of the hypotenuse.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
{
 printf("%f\n", hypot(3.0, 4.0));

 return EXIT_SUCCESS;
}
```

produces the output:

```
5.000000
```

## Classification:

*hypot()* is standard Unix; *hypotf()* is ANSI (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*sqrt*

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(AF_INET,
 SOCK_RAW,
 proto);
```

**Description:**

ICMP is the error- and control-message protocol used by IP and the Internet protocol family. The protocol may be accessed through a “raw socket” for network monitoring and diagnostic functions.

To get the *proto* parameter to *socket()* that’s used to create an ICMP socket, call *getprotobyname()*. You normally use ICMP sockets, which are connectionless, with *sendto()* and *recvfrom()*, although you can also use *connect()* to fix the destination for future packets (in which case you can use the *read()* or *recv()*, and *write()* or *send()* system calls).

Outgoing packets automatically have an IP header prepended to them that’s based on the destination address. Incoming packets are received with the IP header and IP options intact.

**Returns:**

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

**Errors:**

EADDRNOTAVAIL

Tried to create a socket with a network address for which no network interface exists.

EISCONN

Tried to establish a connection on a socket that already has one, or tried to send a datagram with the

destination address specified but the socket is already connected.

ENOBUFS

The system ran out of memory for an internal data structure.

ENOTCONN

Tried to send a datagram, but no destination address was specified, and the socket hasn't been connected.

## See also:

ICMP6, IP protocols

*connect()*, *getprotobyname()*, *read()*, *recv()*, *recvfrom()*, *send()*, *sendto()*, *socket()*, *write()*

*RFC 792*

## Synopsis:

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <netinet/icmp6.h>

int socket(AF_INET6,
 SOCK_RAW,
 proto);
```

## Description:

ICMP6 is the error and control message protocol that IP6 and the Internet Protocol family use. It may be accessed through a “raw socket” for network monitoring and diagnostic functions. Use the *getprotobyname()* function to obtain the *proto* parameter to the *socket()* function, or simply pass IPPROTO\_ICMPV6.

ICMPv6 sockets are connectionless, and are normally used with the *sendto()* and *recvfrom()* functions. You may also use the *connect()* function to fix the destination for future packets (in which case, you may also use the *read()* or *recv()* functions and the *write()* or *send()* system calls).

Outgoing packets automatically have an IP6 header prepended to them (based on the destination address). The ICMP6 pseudo header checksum field (*icmp6\_cksum*, found in the *icmp6\_hdr* structure in *<netinet/icmp6.h>*) is filled automatically by the socket manager. Incoming packets are received without the IP6 header or extension headers.



This behavior is opposite from both IPv4 raw sockets and ICMPv4 sockets.

---

## ICMP6 type/code filter

Each ICMP6 raw socket has an associated filter whose data type is defined as `struct icmp6_filter`. This structure, along with the

macros and constants defined below are defined in the `<netinet/icmp6.h>` header.

You can get and set the current filter by calling `getsockopt()` and `setsockopt()` with a level of `IPPROTO_ICMPV6` and an option name of `ICMP6_FILTER`.

These macros operate on an `icmp6_filter` structure:

```
ICMP6_FILTER_SETPASSALL(struct icmp6_filter *)
ICMP6_FILTER_SETBLOCKALL(struct icmp6_filter *)
ICMP6_FILTER_SETPASS(int, struct icmp6_filter *)
ICMP6_FILTER_SETBLOCK(int, struct icmp6_filter *)
ICMP6_FILTER_WILLPASS(int, const struct icmp6_filter *)
ICMP6_FILTER_WILLBLOCK(int, const struct icmp6_filter *)
```

If the first argument is an integer, it represents an ICMP6 message type, with a value between 0 and 255. The pointer arguments are pointers to the filters that are either set or examined, depending on the macro:

*ICMP6\_FILTER\_SETPASSALL()*

Pass all ICMPv6 messages to the application.

*ICMP6\_FILTER\_SETBLOCKALL()*

Block all ICMPv6 messages from the application.

*ICMP6\_FILTER\_SETPASS()*

Pass messages of a certain ICMPv6 type to the application.

*ICMP6\_FILTER\_SETBLOCK()*

Block messages of a certain ICMPv6 type from the application.

*ICMP6\_FILTER\_WILLPASS()*

Return true or false, depending on whether or not the specified message type is passed to the application.

*ICMP6\_FILTER\_WILLBLOCK()*

Return true or false, depending on whether or not the specified message type is blocked from the application.



When you create an ICMP6 raw socket, it passes all ICMPv6 message types to the application by default.

For more information, see *RFC 2292*.

### See also:

INET6, IP6 protocols

*connect()*, *getprotobyname()*, *getsockopt()*, *read()*, *recv()*, *recvfrom()*,  
*send()*, *sendto()*, *setsockopt()*, *socket()*, *write()*

*RFC 2292*



**Synopsis:**

```
#include <net/if.h>

void if_freenameindex(struct if_nameindex * ptr);
```

**Arguments:**

*ptr* A pointer to the `if_nameindex` structure to be freed.

**Library:**

`libsocket`

**Description:**

The `if_freenameindex()` function frees the dynamic memory that you allocated by calling `if_nameindex()`.

**Classification:**

POSIX 1003.1-2001

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

`getifaddrs()`, `if_indextoname()`, `if_nameindex()`, `if_nametoindex()`



**Synopsis:**

```
#include <net/if.h>

char * if_indextoname(unsigned int ifindex,
 char * ifname);
```

**Arguments:**

*ifindex*    The interface index.

*ifname*    A pointer to a buffer in which *if\_indextoname()* copies the interface name. The buffer must be a minimum of IFNAMSIZ bytes long.

**Library:**

libsocket

**Description:**

The *if\_indextoname()* function maps the interface index specified by *ifindex* to its corresponding name. The name is copied into the buffer pointed to by *ifname*.

**Returns:**

A pointer to the name, or NULL if There isn't an interface corresponding to the specified index.

**Classification:**

POSIX 1003.1-2001

**Safety**

---

Cancellation point    Yes

Interrupt handler    No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*getifaddrs(), if\_freenameindex(), if\_nameindex(), if\_nametoindex()*

## Synopsis:

```
#include <net/if.h>

struct if_nameindex * if_nameindex(void);
```

## Library:

libsocket

## Description:

The *if\_nameindex()* function returns an array of **if\_nameindex** structures, with one structure per interface, as defined in the include file **<net/if.h>**. The **if\_nameindex** structure contains at least the following members:

**unsigned int** *if\_index*

The index of the interface (1, 2, ...).

**char \****if\_name*

A null-terminated name (e.g. **1e0**).

The end of the array of structures is indicated by an entry with an *if\_index* of 0 and an *if\_name* of NULL.

## Returns:

A valid array of **if\_nameindex** structures, or NULL if an error occurred while using *getifaddrs()* to retrieve the list, or there wasn't enough memory available.

## Classification:

POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*getifaddrs(), if\_freenameindex(), if\_indexoname(), if\_nametoindex()*



**Synopsis:**

```
#include <net/if.h>

unsigned int if_nametoindex(const char * ifname);
```

**Arguments:**

*ifname*      The interface name that you want to map.

**Library:**

`libsocket`

**Description:**

The *if\_nametoindex()* function maps the interface name specified by *ifname* to its corresponding index.

**Returns:**

The index number of the interface, or 0 if the specified interface couldn't be found or an error occurred while using *getifaddrs()* to retrieve the list of interfaces.

**Classification:**

POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*getifaddrs(), if\_freenameindex(), if\_indextoname(), if\_nameindex()*

**Synopsis:**

```
#include <ifaddrs.h>

struct ifaddrs {
 struct ifaddrs * ifa_next;
 char * ifa_name;
 u_int ifa_flags;
 struct sockaddr * ifa_addr;
 struct sockaddr * ifa_netmask;
 struct sockaddr * ifa_dstaddr;
 void * ifa_data;
};
```

**Description:**

The **ifaddrs** structure contains the following entries:

|                    |                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ifa_next</i>    | A pointer to the next structure in the list. This field is NULL in the last structure in the list.                                                                                                                                                                             |
| <i>ifa_name</i>    | The interface name.                                                                                                                                                                                                                                                            |
| <i>ifa_flags</i>   | The interface flags, as set by the <b>ifconfig</b> utility.                                                                                                                                                                                                                    |
| <i>ifa_addr</i>    | Either the address of the interface or the link-level address of the interface, if one exists; otherwise it's NULL. See the <i>sa_family</i> member of the <b>sockaddr</b> structure pointed to by <i>ifa_addr</i> to determine the format of the address.                     |
| <i>ifa_netmask</i> | The netmask associated with <b>ifa_addr</b> , if one is set; otherwise it's NULL.                                                                                                                                                                                              |
| <i>ifa_dstaddr</i> | The destination address on a P2P interface, if one exists; otherwise it's NULL. If the interface isn't a P2P interface, <i>ifa_dstaddr</i> contains the broadcast address associated with <i>ifa_addr</i> , if one exists; otherwise it's NULL (see <b>&lt;ifaddr.h&gt;</b> ). |
| <i>ifa_data</i>    | Currently, this is set to NULL.                                                                                                                                                                                                                                                |

**Classification:**

Unix

**See also:**

*freeifaddr(), getifaddr()*

## Synopsis:

```
#include <math.h>

int ilogb (double x);

int ilogbf (float x);
```

## Arguments:

*x* The number you want to compute the integral part of the logarithm.

## Library:

libm

## Description:

The *ilogb()* and *ilogbf()* functions compute the integral part of:

$\log_r |x|$

as a signed integral value, for nonzero finite *x*, where *r* is the radix of the machine's floating point arithmetic.

## Returns:

The exponent part of *x*, in integer format:

| <b>If <i>x</i> is:</b> | <b><i>ilogb()</i> returns:</b> |
|------------------------|--------------------------------|
| 0                      | -INT_MAX                       |
| NAN                    | INT_MAX                        |
| negative infinity      | INT_MAX                        |
| positive infinity      | INT_MAX                        |



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
{
 printf("%f\n", ilogb(.5));

 return EXIT_SUCCESS;
}
```

## Classification:

*ilogb()* is standard Unix; *ilogbf()* is ANSI (draft)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*log()*, *logb()*, *log10()*, *log1p()*

**Synopsis:**

```
#include <hw/inout.h>

uint8_t in8(uintptr_t port);
```

**Arguments:**

*port*     The port you want to read the value from.

**Library:**

`libc`

**Description:**

The *in8()* function reads an 8-bit value from the specified *port*.

**Returns:**

An 8-bit value.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8s()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*, *out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*



## Synopsis:

```
#include <hw/inout.h>

void * in8s(void * buff,
 unsigned len,
 uintptr_t port);
```

## Arguments:

*buff*     A pointer to a buffer where the function can store the values read.

*len*     The number of values that you want to read.

*port*    The port you want to read the values from.

## Library:

`libc`

## Description:

The *in8s()* function reads *len* 8-bit values from the specified *port* and stores them in the buffer pointed to by *buff*.

## Returns:

A pointer to the end of the read data.

## Classification:

QNX Neutrino

### Safety

Cancellation point    No

Interrupt handler     Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*, *out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*

**Synopsis:**

```
#include <hw/inout.h>

uint16_t in16(uintptr_t port);

#define inbe16 (port) ...

#define inle16 (port) ...
```

**Arguments:**

*port*     The port you want to read the value from.

**Library:**

`libc`

**Description:**

The *in16()* function reads a 16-bit value from the specified *port* in native-endian format (there's no conversion required).

The *inbe16()* and *inle16()* macros read a 16-bit value that's in big-endian or little-endian format, respectively, from the specified *port*, and returns the value as native-endian.

**Returns:**

A 16-bit value in native-endian.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    No

Interrupt handler     Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

Both *inbe16()* and *inle16()* are implemented as macros.

**See also:**

*in8()*, *in8s()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*, *out8()*,  
*out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*

**Synopsis:**

```
#include <hw/inout.h>

void * in16s(void * buff,
 unsigned len,
 uintptr_t port);
```

**Arguments:**

*buff*     A pointer to a buffer where the function can store the values read.

*len*       The number of values that you want to read.

*port*      The port you want to read the values from.

**Library:**

libc

**Description:**

The *in16s()* function reads *len* 16-bit values from the specified *port* and stores them in the buffer pointed to by *buff*.

**Returns:**

A pointer to the end of the read data.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point   No

Interrupt handler     Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in8s()*, *in16()*, *in32()*, *in32s()*, *mmap\_device\_io()*, *out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*

**Synopsis:**

```
#include <hw/inout.h>

uint32_t in32(uintptr_t port);

#define inbe32 (port) ...

#define inle32 (port) ...
```

**Arguments:**

*port*     The port you want to read the value from.

**Library:**

`libc`

**Description:**

The *in32()* function reads a 32-bit value from the specified *port*.

The *inbe32()* and *inle32()* macros read a 32-bit value that's in big-endian or little-endian format, respectively, from the specified *port*, and returns the value as native-endian.

**Returns:**

A 32-bit value in native-endian.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    No

Interrupt handler      Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

Both *inbe32()* and *inle32()* are implemented as macros.

**See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32s()*, *mmap\_device\_io()*, *out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*



**Synopsis:**

```
#include <hw/inout.h>

void * in32s(void * buff,
 unsigned len,
 uintptr_t port);
```

**Arguments:**

*buff*     A pointer to a buffer where the function can store the values read.

*len*       The number of values that you want to read.

*port*      The port you want to read the values from.

**Library:**

libc

**Description:**

The *in32s()* function reads *len* 32-bit values from the specified *port* and stores them in the buffer pointed to by *buff*.

**Returns:**

A pointer to the end of the read data.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point   No

Interrupt handler     Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32()*, *mmap\_device\_io()*, *out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*

**Synopsis:**

```
#include <strings.h>

char* index(const char* s,
 int c);
```

**Arguments:**

- s* The string you want to search. This string must end with a null (`\0`) character. The null character is considered to be part of the string.
- c* The character you're looking for.

**Library:**

`libc`

**Description:**

The *index()* function returns a pointer to the first occurrence of the character *c* in the string *s*.

**Returns:**

A pointer to the character, or NULL if the character doesn't occur in the string.

**Classification:**

Legacy Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*rindex(), strchr(), strrchr()*

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

in_addr_t inet_addr(const char * cp);
```

**Arguments:**

*cp*     A pointer to a string that represents an Internet address.

**Library:**

`libsocket`

**Description:**

The *inet\_addr()* routine converts a string representing an IPv4 Internet address (for example, “127.0.0.1”) into a numeric Internet address. To convert a hostname such as `ftp.qnx.com`, call *gethostbyname()*.

All Internet addresses are returned in network byte order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values. For more information on Internet addresses, see *inet\_net\_ntop()*.

**Returns:**

An Internet address, or `INADDR_NONE` if an error occurs.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**Caveats:**

Although the value INADDR\_NONE (0xFFFFFFFF) is a valid broadcast address, *inet\_addr()* always indicates failure when returning that value. The *inet\_aton()* function doesn't share this problem.

**See also:**

*inet\_aton()*, *inet\_network()*

## Synopsis:

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

int inet_aton(const char * cp,
 struct in_addr * addr);
```

## Arguments:

*cp*        A pointer to the character string.

*addr*      A pointer to a `in_addr` structure where the function can store the converted address.

## Library:

`libsocket`

## Description:

The `inet_aton()` routine interprets the specified character string as an IPv4 Internet address, placing the address into the structure provided.

All Internet addresses are returned in network byte order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values.

For more information on Internet addresses, see `inet_net_ntop()`.

## Returns:

- 1        Success; the string was successfully interpreted.
- 0        Failure; the string is invalid.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*gethostbyname(), getnetent() inet\_addr(), inet\_lnaof(),  
inet\_makeaddr(), inet\_netof(), inet\_network(), inet\_ntoa()  
/etc/hosts, /etc/networks in the Utilities Reference*



**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

unsigned long inet_lnaof(struct in_addr in);
```

**Arguments:**

*in*     An Internet address.

**Library:**

`libsocket`

**Description:**

The *inet\_lnaof()* routine returns the local network address for an IPv4 Internet address.

All Internet addresses are returned in network byte order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values. For more information on Internet addresses, see *inet\_net\_ntop()*.

**Returns:**

A local network address.

**Classification:**

Standard Unix, POSIX 1003.1g (draft)

**Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*inet\_aton(), inet\_netof()*

*Convert a network number and a local network address into an Internet address*

## Synopsis:

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

struct in_addr inet_makeaddr(unsigned long net,
 unsigned long lna);
```

## Arguments:

*net*     An Internet network number.

*lna*     The local network address.

## Library:

`libsocket`

## Description:

The *inet\_makeaddr()* routine takes an Internet network number and a local network address and constructs an IPv4 Internet address.

All Internet addresses are returned in network byte order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values. For more information on Internet addresses, see *inet\_net\_ntop()*.

## Returns:

An Internet address.

## Classification:

Standard Unix, POSIX 1003.1g (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*inet\_aton()*

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

char * inet_net_ntop(int af,
 const void * src,
 int bits,
 char * dst,
 size_t size);
```

**Arguments:**

- af* The address family. Currently, only AF\_INET is supported.
- src* A pointer to the Internet network number that you want to convert. The format of the address is interpreted according to *af*.
- bits* The number of bits that specify the network number (*src*).
- dst* a pointer to the buffer where the function can store the converted address.
- size* The size of the buffer that *dst* points to, in bytes.

**Library:**

`libsocket`

**Description:**

The *inet\_net\_ntop()* function converts an Internet network number from network format (usually a `struct in_addr` or some other binary form, in network byte order) to CIDR (Classless Internet Domain Routing) presentation format that's suitable for external display purposes.

With CIDR, a single IP address can be used to designate many unique IP addresses. A CIDR IP address looks like a normal IP address,

except that it ends with a slash (/) followed by a number, called the *IP prefix*. For example:

`172.200.0.0/16`

The IP prefix specifies how many addresses are covered by the CIDR address, with lower numbers covering more addresses.

### Network Numbers (IPv4 Internet addresses)

You can specify Internet addresses in the “dotted quad” notation, or Internet network numbers, using one of the following forms:

*a.b.c.d/bits* or *a.b.c.d*

When you specify a four-part address, each part is interpreted as a byte of data and is assigned, from left to right, to the four bytes of an Internet network number (or Internet address). When an Internet network number is viewed as a 32-bit integer quantity on a system that uses little-endian byte order (i.e. right to left), such as the Intel 386, 486 and Pentium processors, the bytes referred to above appear as “*d.c.b.a*”.

*a.b.c* When you specify a three-part address, the last part is interpreted as a 16-bit quantity and is placed in the rightmost two bytes of the Internet network number (or network address). This makes the three-part address format convenient for specifying Class B network addresses as *net.net.host*.

*a.b* When you specify a two-part address, the last part is interpreted as a 24-bit quantity and is placed in the rightmost three bytes of the Internet network number (or network address). This makes the two-part number format convenient for specifying Class A network numbers as *net.host*.

*a* When you specify a one-part address, the value is stored directly in the Internet network number (network address) without any byte rearrangement.

All numbers supplied as “parts” in a dot notation may be decimal, octal, or hexadecimal, as specified in the C language. That is, a number is interpreted as decimal unless it has a leading 0 (octal), or a leading 0x or 0X (hex).

**Returns:**

A pointer to the destination string (*dst*), or NULL if a system error occurs (*errno* is set).

**Errors:**

ENOENT      Invalid argument *af*.

**Classification:**

POSIX 1003.1-2000

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*inet\_aton()*, *inet\_net\_ntop()*

## ***inet\_netof()***

© 2004, QNX Software Systems Ltd.

*Extract the network number from an Internet address*

### **Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

unsigned long inet_netof(struct in_addr in);
```

### **Arguments:**

*in*     An Internet address.

### **Library:**

`libsocket`

### **Description:**

The *inet\_netof()* routine returns the network number of the specified IPv4 Internet address.

All Internet addresses are returned in network order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values. For more information on Internet addresses, see *inet\_net\_ntop()*.

### **Returns:**

An Internet network number.

### **Classification:**

Standard Unix, POSIX 1003.1g (draft)

#### **Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*



**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*inet\_aton()*, *inet\_lnaof()*

## ***inet\_net\_pton()***

© 2004, QNX Software Systems Ltd.

*Convert an Internet network number from CIDR format to network format*

### **Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

int inet_net_pton(int af,
 const char * src,
 void * dst,
 size_t size);
```

### **Arguments:**

- af*        The address family. Currently, only AF\_INET is supported.
- src*        A pointer to the presentation-format (CIDR) address. The format of the address is interpreted according to *af*.
- dst*        A pointer to the buffer where the function can store the converted address.
- size*       The size of the buffer pointed to by *dst*, in bytes.

### **Library:**

`libsocket`

### **Description:**

The *inet\_net\_pton()* function converts an Internet network number from presentation format — a printable form as held in a character string, such as, Internet standard dot notation, or Classless Internet Domain Routing (CIDR) — to network format (usually a struct `in_addr` or some other internal binary representation, in network byte order).

For more information on Internet addresses, see *inet\_net\_ntop()*.

## Returns:

The number of bits that specify the network number (computed based on the class, or specified with /CIDR), or -1 if an error occurred (*errno* is set).

## Errors:

ENOENT      Invalid argument *af*.

## Classification:

POSIX 1003.1-2000

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*inet\_aton()*, *inet\_net\_ntop()*

## ***inet\_network()***

© 2004, QNX Software Systems Ltd.

*Convert a string into an Internet network number*

### **Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

unsigned long inet_network(const char * cp);
```

### **Arguments:**

*cp*     A pointer to a string representing an Internet address.

### **Library:**

`libsocket`

### **Description:**

The *inet\_network()* routine converts a string representing an IPv4 Internet address (for example, “127.0.0.1”) into a numeric Internet network number.

All Internet addresses are returned in network order (bytes are ordered from left to right). All network numbers and local address parts are returned as machine-format integer values. For more information on Internet addresses, see *inet\_net\_ntop()*.

### **Returns:**

An Internet network number, or INADDR\_NONE if an error occurs.

### **Classification:**

Standard Unix, POSIX 1003.1g (draft)

#### **Safety**

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*inet\_addr()*, *inet\_aton()*

## ***inet\_ntoa()***

© 2004, QNX Software Systems Ltd.

*Convert an Internet address into a string*

### **Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

char * inet_ntoa(struct in_addr in);
```

### **Arguments:**

*in*     The Internet address that you want to convert.

### **Library:**

`libsocket`

### **Description:**

The *inet\_ntoa()* routine converts an IPv4 Internet address into an ASCII string representing the address in dot notation (for example, "127.0.0.1").

For more information on Internet addresses, see *inet\_net\_ntop()*.

### **Returns:**

A string representing an Internet address.

### **Classification:**

Standard Unix, POSIX 1003.1-2001

#### **Safety**

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**Caveats:**

The string returned by this function is stored in a static buffer that's reused for every call to *inet\_ntoa()*. For a thread-safe version, see *inet\_ntoa\_r()*.

**See also:**

*inet\_aton()*, *inet\_ntoa\_r()*

## ***inet\_ntoa\_r()***

© 2004, QNX Software Systems Ltd.

*Convert an Internet address into a string*

### **Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

char * inet_ntoa_r(struct in_addr in,
 char * buffer,
 int buflen);
```

### **Arguments:**

*in*            The Internet address that you want to convert.

*buffer*        A buffer where the function can store the result.

*buflen*        The size of the buffer, in bytes.

### **Library:**

libsocket

### **Description:**

The *inet\_ntoa\_r()* function is a thread-safe version of *inet\_ntoa()*. It converts an IPv4 Internet address into a string (for example, "127.0.0.1"). For more information on this routine, see *inet\_aton()*.

### **Returns:**

A string representing an Internet address, or NULL if an error occurs (*errno* is set).

### **Errors:**

ERANGE        The supplied *buffer* isn't large enough to store the result.



## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*inet\_aton()*, *inet\_ntoa()*

## ***inet\_ntop()***

© 2004, QNX Software Systems Ltd.

*Convert a numeric network address to a string*

### **Synopsis:**

```
#include <sys/socket.h>
#include <arpa/inet.h>

const char * inet_ntop(int af,
 const void * src,
 char * dst,
 socklen_t size);
```

### **Arguments:**

- af*      The *src* address's network family; one of:
- AF\_INET    IPv4 addresses
  - AF\_INET6   IPv6 addresses
- src*      The numeric network address that you want to convert to a string.
- dst*      The text string that represents the translated network address. You can use the following constants to allocate buffers of the correct size (they're defined in `<netinet/in.h>`):
- INET\_ADDRSTRLEN — storage for an IPv4 address
  - INET6\_ADDRSTRLEN — storage for an IPv6 address
- size*     The size of the buffer pointed to by *dst*.

### **Library:**

libc

### **Description:**

The *inet\_ntop()* function converts a numeric network address pointed to by *src* into a text string in the buffer pointed to by *dst*.

## Returns:

A pointer to the buffer containing the text version of the address, or NULL if an error occurs (*errno* is set).

## Errors:

EAFNOSUPPORT

The value of the *af* argument isn't a supported network family.

ENOSPC

The *dst* buffer isn't large enough (according to *size*) to store the translated address.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/socket.h>
#include <arpa/inet.h>
#include <errno.h>

#define INADDR "10.1.0.29"
#define IN6ADDR "DEAD:BEEF:7654:3210:FEDC:3210:7654:BA98"

int
main()
{
 struct in_addr inaddr;
 struct in6_addr in6addr;
 char buf[INET_ADDRSTRLEN], buf6[INET6_ADDRSTRLEN];
 int rval;

 if ((rval = inet_pton(AF_INET, INADDR, &inaddr)) == 0) {
 printf("Invalid address: %s\n", INADDR);
 exit(EXIT_FAILURE);
 } else if (rval == -1) {
 perror("inet_pton");
 exit(EXIT_FAILURE);
 }

 if (inet_ntop(AF_INET, &inaddr, buf, sizeof(buf)) != NULL)
 printf("inet addr: %s\n", buf);
 else {
 perror("inet_ntop");
 exit(EXIT_FAILURE);
 }
}
```

```
if ((rval = inet_pton(AF_INET6, IN6ADDR, &in6addr)) == 0) {
 printf("Invalid address: %s\n", IN6ADDR);
 exit(EXIT_FAILURE);
} else if (rval == -1) {
 perror("inet_pton");
 exit(EXIT_FAILURE);
}

if (inet_ntop(AF_INET6, &in6addr, buf6, sizeof(buf6)) != NULL)
 printf("inet6 addr: %s\n", buf6);
else {
 perror("inet_ntop");
 exit(EXIT_FAILURE);
}

return(EXIT_SUCCESS);
}
```

**Classification:**

Unix, POSIX 1003.1-2001

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*inet\_pton()*

## Synopsis:

```
#include <sys/socket.h>
#include <arpa/inet.h>

int inet_pton(int af,
 const char * src,
 void * dst);
```

## Arguments:

- af*      The *src* address's network family; one of:
- AF\_INET    IPv4 addresses
  - AF\_INET6   IPv6 addresses
- src*      A pointer to the text host address that you want to convert. The format of the address is interpreted according to *af*
- dst*      A pointer to a buffer where the function can store the converted address.

## Library:

`libc`

## Description:

The *inet\_pton()* function converts the standard text representation of the numeric network address (*src*) into its numeric network byte-order binary form (*dst*).

The converted address is stored in network byte order in *dst*. The buffer pointed to by *dst* must be large enough to hold the numeric address:

| Family   | Numeric address size |
|----------|----------------------|
| AF_INET  | 4 bytes              |
| AF_INET6 | 16 bytes             |

## AF\_INET addresses

IPv4 addresses must be specified in the standard dotted-decimal form:

*ddd.ddd.ddd.ddd*

where *ddd* is a one- to three-digit decimal number between 0 and 255.



---

Many existing implementations of *inet\_addr()* and *inet\_aton()* accept nonstandard input: octal numbers, hexadecimal numbers, and fewer than four numbers. The *inet\_pton()* function doesn't accept these formats.

---

## AF\_INET6 addresses

IPv6 addresses must be specified in one of the following standard formats:

- The preferred form is:

*xxxx:xxxx:xxxx*

where *x* is a hexadecimal value for one of the eight 16-bit pieces of the address. For example:

**DEAD:BEEF:7654:3210:FEDC:3210:7654:BA98**  
**417A:200C:800:8:0:0:0:1080**

- A **::** can be used once per address to represent multiple groups of 16 zero-bits. For example, the following addresses:

**1080:0:0:0:8:800:200C:417A**  
**FF01:0:0:0:0:0:0:43**  
**0:0:0:0:0:0:0:1**  
**0:0:0:0:0:0:0:0**

can be represented as:

```
1080::8:800:200C:417A
FF01::43
::1
::
```

- A convenient format when dealing with mixed IPv4 and IPv6 environments is:

```
x:x:x:x:x:d.d.d.d
```

where *x* is a hexadecimal value for one of the six high-order 16-bit pieces of the address and *d* is a decimal value for one of the four low-order 8-bit pieces of the address (standard AF\_INET representation). For example:

```
0:0:0:0:0:0:13.1.68.3
0:0:0:0:0:FFFF:129.144.52.38
Or, in their compressed forms:
::13.1.68.3
::FFFF:129.144.52.38
```

## Returns:

- 1 Success.
- 0 The input isn't a valid address.
- 1 An error occurred (*errno* is set).

## Errors:

EAFNOSUPPORT

The *af* argument isn't one of the supported networking families.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/socket.h>
#include <arpa/inet.h>
#include <errno.h>

#define INADDR "10.1.0.29"
```

```
#define IN6ADDR "DEAD:BEEF:7654:3210:FEDC:3210:7654:BA98"

int
main()
{
 struct in_addr inaddr;
 struct in6_addr in6addr;
 char buf[INET_ADDRSTRLEN], buf6[INET6_ADDRSTRLEN];
 int rval;

 if ((rval = inet_pton(AF_INET, INADDR, &inaddr)) == 0) {
 printf("Invalid address: %s\n", INADDR);
 exit(EXIT_FAILURE);
 } else if (rval == -1) {
 perror("inet_pton");
 exit(EXIT_FAILURE);
 }

 if (inet_ntop(AF_INET, &inaddr, buf, sizeof(buf)) != NULL)
 printf("inet addr: %s\n", buf);
 else {
 perror("inet_ntop");
 exit(EXIT_FAILURE);
 }

 if ((rval = inet_pton(AF_INET6, IN6ADDR, &in6addr)) == 0) {
 printf("Invalid address: %s\n", IN6ADDR);
 exit(EXIT_FAILURE);
 } else if (rval == -1) {
 perror("inet_pton");
 exit(EXIT_FAILURE);
 }

 if (inet_ntop(AF_INET6, &in6addr, buf6, sizeof(buf6)) != NULL)
 printf("inet6 addr: %s\n", buf6);
 else {
 perror("inet_ntop");
 exit(EXIT_FAILURE);
 }

 return(EXIT_SUCCESS);
}
```



## Classification:

Unix, POSIX 1003.1-2001

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*inet\_ntop()*

*RFC 2373*

## ***inet6\_option\_\****(*)*

© 2004, QNX Software Systems Ltd.

*Manipulate IPv6 hop-by-hop and destination options*

### **Synopsis:**

```
#include <netinet/in.h>

int inet6_option_space(int nbytes);

int inet6_option_init(void *bp,
 struct cmsghdr **cmsgp,
 int type);

int inet6_option_append(struct cmsghdr *cmsg,
 const u_int8_t *typep,
 int multx,
 int plusy);

u_int8_t * inet6_option_alloc(struct cmsghdr *cmsg,
 int datalen,
 int multx,
 int plusy);

int inet6_option_next(const struct cmsghdr *cmsg,
 u_int8_t **tptrp);

int inet6_option_find(const struct cmsghdr *cmsg,
 u_int8_t **tptrp,
 int type);
```

### **Arguments:**

|               |                                                                                                                                                                                                                                           |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>nbytes</i> | Size of the structure that defines the option. It includes any pad bytes at the beginning the value <i>y</i> in the alignment term $xn + y$ , the type byte, the length byte, and the option data.                                        |
| <i>bp</i>     | Pointer to previously allocated space that contains the ancillary data object. It must be large enough to contain all the individual options to be added by later calls to <i>inet6_option_append()</i> and <i>inet6_option_alloc()</i> . |
| <i>cmsgp</i>  | Pointer to a <b>cmsghdr</b> structure. The <i>*cmsgp</i> variable is initialized by this function to point to the <b>cmsghdr</b>                                                                                                          |

structure constructed by this function in the buffer pointed to by *bp*.

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>  | Either IPV6_HOPOPTS or IPV6_DSTOPTS. This type is stored in the <i>msg_type</i> member of the <b>cmsghdr</b> structure pointed to by <i>*msgp</i> .                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <i>msg</i>   | Pointer to the <b>cmsghdr</b> structure that must have been initialized by <i>inet6_option_init()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <i>typep</i> | Pointer to the 8-bit option type. It's assumed that this field is immediately followed by the 8-bit option data length field, which is then followed by the option data. The caller initializes these three fields (the type-length-value, or TLV) before calling this function.<br><br>The option type must have a value from 2 to 255, inclusive. (0 and 1 are reserved for the <b>Pad1</b> and <b>PadN</b> options, respectively.)<br><br>The option data length must have a value between 0 and 255, inclusive, and is the length of the option data that follows. |
| <i>multx</i> | value <i>x</i> in the alignment term $xn + y$ . It must have a value of 1, 2, 4, or 8.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>plusy</i> | Value <i>y</i> in the alignment term $xn + y$ . It must have a value between 0 and 7, inclusive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>tptrp</i> | Pointer to a pointer to an 8-bit byte.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

**Library:****libsocket****Description:**

These functions perform hop-by-hop and destination options following *RFC2292* that alleviates alignment constraints, padding and ancillary data manipulation. You can find the prototypes for these functions in the `<netinet/in.h>` header.

## inet6\_option\_space()

This function returns the number of bytes required to hold an option when it's stored as ancillary data, including the `cmsghdr` structure at the beginning, and any padding at the end (to make its size a multiple of 8 bytes). The argument is the size of the structure that defines the option. It includes any pad bytes at the beginning (the value  $y$  in the alignment term  $xn + y$ ), the type byte, the length byte, and the option data.



---

When multiple options are stored in a single ancillary data object, this function overestimates the amount of space required by the size of  $N-1$  `cmsghdr` structures, where  $N$  is the number of options to be stored in the object. This is of little consequence, since it's assumed that most hop-by-hop option and destination option headers carry only one option (see Appendix B of *RFC 2460*).

---

## inet6\_option\_init()

This function is called once per ancillary data object that contains either hop-by-hop or destination options.

- |              |                                                                                                                                                                                                                                                       |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>bp</i>    | Pointer to previously allocated space that contains the ancillary data object. It must be large enough to contain all the individual options to be added by later calls to <code>inet6_option_append()</code> and <code>inet6_option_alloc()</code> . |
| <i>cmsgp</i> | Pointer to a <code>cmsghdr</code> structure. The <code>*cmsgp</code> variable is initialized by this function to point to the <code>cmsghdr</code> structure constructed by this function in the buffer pointed to by <i>bp</i> .                     |
| <i>type</i>  | Either <code>IPV6_HOPOPTS</code> or <code>IPV6_DSTOPTS</code> . This type is stored in the <code>cmsg_type</code> member of the <code>cmsghdr</code> structure pointed to by <code>*cmsgp</code> .                                                    |

### Returns:

- 0 Success.
- 1 An error has occurred.

***inet6\_option\_append()***

This function appends a hop-by-hop option or a destination option into an ancillary data object that has been initialized by *inet6\_option\_init()*.

- msg* Pointer to the **cmsghdr** structure that must have been initialized by *inet6\_option\_init()*.
- type* Pointer to the 8-bit option type. It's assumed that this field is immediately followed by the 8-bit option data length field, which is then followed by the option data. The caller initializes these three fields (the type-length-value, or TLV) before calling this function.  
  
The option type must have a value from 2 to 255, inclusive. (0 and 1 are reserved for the **Pad1** and **PadN** options, respectively.)  
  
The option data length must have a value between 0 and 255, inclusive, and is the length of the option data that follows.
- multx* value *x* in the alignment term  $xn + y$ . It must have a value of 1, 2, 4, or 8.
- plusy* Value *y* in the alignment term  $xn + y$ . It must have a value between 0 and 7, inclusive.

**Returns:**

- 0 Success.
- 1 An error has occurred.

***inet6\_option\_alloc()***

This function appends a hop-by-hop option or a destination option into an ancillary data object that has been initialized by *inet6\_option\_init()*.

The difference between this function and *inet6\_option\_append()* is that the latter copies the contents of the previously built option into the ancillary data object. This function returns a pointer to the space in the data object where the option's TLV must then be built by the caller.

|                |                                                                                                                                                                                                                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>msg</i>     | pointer to the <b>cmsghdr</b> structure that must have been initialized by <i>inet6_option_init()</i> .                                                                                                                                                                                                                                   |
| <i>datalen</i> | Value of the option data length byte for this option. This value is required as an argument to allow the function to determine if padding should be appended at the end of the option, argument since the option data length must already be stored by the caller. (The <i>inet6_option_append()</i> function doesn't need a data length) |
| <i>multx</i>   | Value <i>x</i> in the alignment term $xn + y$ . It must have a value of 1, 2, 4, or 8.                                                                                                                                                                                                                                                    |
| <i>plusy</i>   | Value of <i>y</i> in the alignment term $xn + y$ . It must have a value between 0 and 7, inclusive.                                                                                                                                                                                                                                       |

**Returns**

Pointer to the 8-bit option type field that starts the option

Success.

NULL An error has occurred.

***inet6\_option\_next()***

This function processes the next hop-by-hop option or destination option in an ancillary data object. If another option remains to be processed, the return value of the function is 0 and *\*tprp* points to the 8-bit option type field the option data.

The *cmsg* variable is a pointer to **cmsghdr** structure for which *cmsg\_Level* equals IPPROTO\_IPV6 and *cmsg\_type* equals either IPV6\_HOPOPTS or IPV6\_DSTOPTS.

The *tprp* is a pointer to a pointer to an 8-bit byte and *\*tprp* is used by the function to remember its place in the ancillary data object each time the function is called. The first time this function is called for a given ancillary data object, *\*tprp* must be set to NULL.

Each time this function returns success, *\*tprp* points to the 8-bit option type field for the next option to be processed.

**Returns:**

- 0 The option is located and the *\*tprp* points to the 8-bit option type field.
- 1 with *\*tprp* pointing to NULL  
No more options to process.
- 1 with *\*tprp* pointing to non-NULL  
An error has occurred.

***inet6\_option\_find()***

This function is similar to *inet6\_option\_next()*. It however, lets the caller specify the option type to be searched for, instead of always returning the next option in the ancillary data object. The *cmsg* is a pointer to the **cmsghdr** structure of which *cmsg\_Level* equals IPPROTO\_IPV6 and *cmsg\_type* equals either IPV6\_HOPOPTS or IPV6\_DSTOPTS.

The *tprp* is a pointer to a pointer to an 8-bit byte that is used by the function to remember its place in the ancillary data object each time the function is called.

The first time this function is called for a given ancillary data object, *\*tprp* must be set to NULL. This function starts searching for an option of the specified type beginning after the value of *\*tprp* pointer.

**Returns:**

0 with *\*tprp* pointing to the 8-bit option

The option is located.

-1 with *\*tprp* pointing to NULL

The option is not located.

-1 with *\*tprp* pointing to non-NULL

An error has occurred.

**Classification:**

Standard *RFC 2292*

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

W. Stevens and M. Thomas, Advanced Sockets API for IPv6, *RFC 2292*, February 1998. Contains examples.

S. Deering and R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, *RFC 2460*, December 1998.



**Synopsis:**

```
#include <netinet/in.h>

struct sockaddr_in6 {
 uint8_t sin6_len;
 sa_family_t sin6_family;
 in_port_t sin6_port;
 uint32_t sin6_flowinfo;
 struct in6_addr sin6_addr;
 uint32_t sin6_scope_id;
};
```

**Description:****Protocols**

The INET6 family consists of the:

- IPv6 network protocol
- Internet Control Message Protocol version 6 (ICMP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP).

TCP supports the SOCK\_STREAM abstraction, while UDP supports the SOCK\_DGRAM abstraction. Note that TCP and UDP are common to INET and INET6. A raw interface to IPv6 is available by creating an Internet SOCK\_RAW socket. The ICMPv6 message protocol may be accessed from a raw socket.

The INET6 protocol family is an updated version of the INET family. While INET implements Internet Protocol version 4, INET6 implements Internet Protocol version 6.

**Addressing**

IPv6 addresses are 16-byte quantities, stored in network standard (big-endian) byte order. The header file `<netinet/in.h>` defines this address as a discriminated union.

Sockets bound to the INET6 family use the structure shown above.

You can create sockets with the local address `::` (which is equal to IPv6 address `0:0:0:0:0:0:0:0`) to cause “wildcard” matching on incoming messages. You can specify the address in a call to `connect()` or `sendto()` as `::` to mean the local host. You can get the `::` value by setting the `sin6_addr` field to 0, or by using the address contained in the `in6addr_any` global variable, which is declared in `<netinet6/in6.h>`.

The IPv6 specification defines *scoped addresses*, such as link-local or site-local addresses. A scoped address is ambiguous to the kernel if it’s specified without a scope identifier. To manipulate scoped addresses properly in your application, use the advanced API defined in *RFC 2292*. A compact description on the advanced API is available in IP6. If you specify scoped addresses without an explicit scope, the socket manager may return an error.



---

Scoped addresses are currently experimental, from both a specification and an implementation point of view.

---

The KAME implementation supports extended numeric IPv6 address notation for link-local addresses. For example, you can use `fe80::1%de0` to specify “`fe80::1` on the `de0` interface.” The `getaddrinfo()` and `getnameinfo()` functions support this notation. Some utilities, such as `telnet` and `ftp`, can use the notation. With special programs like `ping6`, you can disambiguate scoped addresses by specifying the outgoing interface with extra command-line options.

The socket manager handles scoped addresses in a special manner. In the socket manager’s routing tables or interface structures, a scoped address’s interface index is embedded in the address. Therefore, the address contained in some of the socket manager structures isn’t the same as on the wire. The embedded index becomes visible when using the `PF_ROUTE` socket or the `sysctl()` function. You shouldn’t use the embedded form.

## Interaction between IPv4/v6 sockets

The behavior of the AF\_INET6 TCP/UDP socket is documented in the *RFC 2553* specification, which states:

- A specific bind on an AF\_INET6 socket (*bind()* with an address specified) should accept IPv6 traffic to that address only.
- If you perform a wildcard bind on an AF\_INET6 socket (*bind()* to the IPv6 address `:::`), and there isn't a wildcard-bound AF\_INET socket on that TCP/UDP port, then the IPv6 traffic as well as the IPv4 traffic should be routed to that AF\_INET6 socket. IPv4 traffic should be seen by the application as if it came from an IPv6 address such as `::ffff:10.1.1.1`. This is called an *IPv4 mapped address*.
- If there are both wildcard-bound AF\_INET sockets and wildcard-bound AF\_INET6 sockets on one TCP/UDP port, they should operate independently: IPv4 traffic should be routed to the AF\_INET socket, and IPv6 should be routed to the AF\_INET6 socket.

However, the *RFC 2553* specification doesn't define the constraint between the binding order, nor how the IPv4 TCP/UDP port numbers and the IPv6 TCP/UDP port numbers relate each other (whether they must be integrated or separated). The behavior is very different from implementation to implementation. It is unwise to rely too much on the behavior of the AF\_INET6 wildcard-bound socket. Instead, connect to two sockets, one for AF\_INET and another for AF\_INET6, when you want to accept both IPv4 and IPv6 traffic.



---

**CAUTION:** Use caution when handling connections from IPv4 mapped addresses with AF\_INET6 sockets — if the target node routes IPv4 traffic to AF\_INET6 sockets, malicious parties can bypass security.

---

Because of the security hole, by default, NetBSD doesn't route IPv4 traffic to AF\_INET6 sockets. If you want to accept both IPv4 and IPv6 traffic, use two sockets. IPv4 traffic may be routed with multiple

per-socket/per-node configurations, but, it isn't recommended. See IP6 for details.



---

The IPv6 support is subject to change as the Internet protocols develop. Don't depend on details of the current implementation, but rather the services exported. Try to implement version-independent code as much as possible, because you'll need to support both INET and INET6.

---

### See also:

ICMP, ICMP6, IP6, IP, TCP, UDP protocols

*bind()*, *connect()*, *getaddrinfo()*, *ioctl()*, *sendto()*, *socket()*, *sysctl()*

**ftp**, **ping6**, **telnet** in the *Utilities Reference*

*RFC 2553*, *RFC 2292*

**Synopsis:**

```
#include <netinet/in.h>
size_t inet6_rthdr_space(int type,
 int segments);

struct cmsghdr * inet6_rthdr_init(void *bp,
 int type);

int inet6_rthdr_add(struct cmsghdr *cmsg,
 const struct in6_addr *addr,
 unsigned int flags);

int inet6_rthdr_lasthop(struct cmsghdr *cmsg,
 unsigned int flags);

int inet6_rthdr_reverse(const struct cmsghdr *in,
 struct cmsghdr *out);

int inet6_rthdr_segments(const struct cmsghdr *cmsg);

struct in6_addr * inet6_rthdr_getaddr(struct cmsghdr *cmsg,
 int index);

int inet6_rthdr_getflags(const struct cmsghdr *cmsg,
 int index);
```

**Arguments:**

|                 |                                                                                                                    |
|-----------------|--------------------------------------------------------------------------------------------------------------------|
| <i>type</i>     | The type of IPv6 Routing header (e.g. Type 0 as defined in <netinet/in.h>).                                        |
| <i>segments</i> | The number of segments (addresses) in the Routing header.                                                          |
| <i>bp</i>       | Pointer to the buffer that contains a <b>cmsghdr</b> structure followed by a Routing header of the specified type. |
| <i>addr</i>     | IPv6 address structure.                                                                                            |
| <i>flags</i>    | Routing header flags.                                                                                              |

|              |                                                                              |
|--------------|------------------------------------------------------------------------------|
| <i>in</i>    | Ancillary data containing Routing header.                                    |
| <i>out</i>   | Ancillary data containing Routing header.                                    |
| <i>cmsg</i>  | Ancillary data containing Routing header.                                    |
| <i>index</i> | A value between 0 and the number returned by <i>inet6_rthdr_segments()</i> . |

**Library:**

`libsocket`

**Description:**

Your application can now call eight functions to build and examine a Routing header. The function prototypes for these functions are all in the `<netinet/in.h>` header.

The following functions build a Routing header:

| <b>Use this function:</b>    | <b>To:</b>                                         |
|------------------------------|----------------------------------------------------|
| <i>inet6_rthdr_space()</i>   | Return number of bytes required for ancillary data |
| <i>inet6_rthdr_init()</i>    | Initialize ancillary data for Routing header       |
| <i>inet6_rthdr_add()</i>     | Add IPv6 address and flags to Routing header       |
| <i>inet6_rthdr_lasthop()</i> | Specify the flags for the final hop                |

The following functions deal with a returned Routing header:

| <b>Use this function:</b>     | <b>To:</b>                                    |
|-------------------------------|-----------------------------------------------|
| <i>inet6_rthdr_reverse()</i>  | Reverse a Routing header                      |
| <i>inet6_rthdr_segments()</i> | Return number of segments in a Routing header |
| <i>inet6_rthdr_getaddr()</i>  | Fetch one address from a Routing header       |
| <i>inet6_rthdr_getflags()</i> | Fetch one flag from a Routing header          |

***inet6\_rthdr\_space()***

This function returns the number of bytes required to hold a Routing header of the specified type containing a specified number of segments (addresses). For an IPv6 Type 0 Routing header, the number of segments must be between 1 and 23, inclusive. The return value includes the size of the **cmsghdr** structure that precedes the Routing header, and any required padding.

If the return value is 0, then either the type of the Routing header isn't supported by this implementation or the number of segments is invalid for this type of Routing header.



This function returns the size but doesn't allocate the space required for the ancillary data. This allows an application to allocate a larger buffer, if other ancillary data objects are desired. All the ancillary data objects must be specified to *sendmsg()* as a single **msg\_control** buffer in the **msghdr** structure *msg\_control* member.

***inet6\_rthdr\_init()***

This function initializes the buffer pointed to by *bp* to contain a **cmsghdr** structure followed by a Routing header of the specified type. The *cmsghdr* member of the **cmsghdr** structure is initialized to the size of the structure plus the amount of space required by the Routing header.

The *msg\_level* and *msg\_type* members are also initialized as required.

The caller must allocate the buffer and its size that is determined by calling *inet6\_rthdr\_space()*.

Upon success, the return value is the pointer to the **cmsghdr** structure, and this is then used as the first argument to the next two functions.

The function returns NULL on error.

### ***inet6\_rthdr\_add()***

This function adds the address pointed to by *addr* to the end of the Routing header being constructed and sets the type of this hop to the value of *flags*. For an IPv6 Type 0 Routing header, *flags* must be either IPV6\_RTHDR\_LOOSE or IPV6\_RTHDR\_STRICT.

If successful, the *msg\_len* member of the **cmsghdr** structure is updated to account for the new address in the Routing header.

#### **Returns:**

- 0 Success.
- 1 An error has occurred.

### ***inet6\_rthdr\_lasthop()***


This function specifies the Strict/Loose flag for the final hop of a Routing header. For an IPv6 Type 0 Routing header, *flags* must be either IPV6\_RTHDR\_LOOSE or IPV6\_RTHDR\_STRICT.

#### **Returns:**

- 0 Success.
- 1 An error has occurred.




---

 A Routing header specifying N intermediate nodes requires N+1 Strict/Loose flags. This requires N calls to *inet6\_rthdr\_add()* followed by one call to *inet6\_rthdr\_lasthop()*.

---

### ***inet6\_rthdr\_reverse()***

 The *inet6\_rthdr\_reverse()* has not been implemented yet.

---

This function takes a Routing header that has been received as ancillary data (pointed to by the first argument, *in*) and writes a new Routing header. The Routing header sends datagrams along the reverse of that route. Both arguments are allowed to point to the same buffer (that is, the reversal can occur in place).

#### **Returns:**

0      Success.  
-1     An error has occurred.

### ***inet6\_rthdr\_segments()***

This function returns the number of segments (addresses) contained in the Routing header described by *msg*.

#### **Returns:**

1 to 23    Success.  
-1         An error has occurred.

### ***inet6\_rthdr\_getaddr()***

This function returns a pointer to the **IPv6** address specified by *index* in the Routing header described by *msg*. The *index* must have a value between 1 and the number returned by *inet6\_rthdr\_segments()*. An application should first call *inet6\_rthdr\_segments()* to obtain the number of segments in the Routing header.

The function returns NULL on error.

***inet6\_rthdr\_getflags()***

This function returns the flags value specified by *index* in the Routing header described by *cmsg*. The *index* must have a value between 0 and the number returned by *inet6\_rthdr\_segments()*. For an IPv6 Type 0 Routing header, the return value is either IPV6\_RTHDR\_LOOSE or IPV6\_RTHDR\_STRICT.

The function returns -1 on error.



---

Addresses are indexed starting at 1, and flags starting at 0. They're consistent with the terminology and figures in *RFC2460*.

---

**Classification:**

*RFC 2292*

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

W. Stevens and M. Thomas, Advanced Sockets API for IPv6, *RFC 2292*, February 1998. Contains good examples.

S. Deering and R. Hinden, Internet Protocol, Version 6 (IPv6) Specification, *RFC 2460*, December 1998.

### Synopsis:

```
#include <grp.h>
#include <sys/types.h>

int initgroups(const char * name,
 gid_t basegid);
```

### Arguments:

|                |                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------|
| <i>name</i>    | The name of the user whose group membership you want to use as the supplementary group access list. |
| <i>basegid</i> | A group ID that you want to include in the group access list.                                       |

### Library:

libc

### Description:

The *initgroups()* function reads the group membership for the user specified by *name* from the group database, and then initializes the supplementary group access list of the calling process (see *getgrnam()* and *getgroups()*).

If the number of groups in the supplementary access list exceeds NGROUPS\_MAX, the extra groups are ignored.

### Returns:

|    |                                           |
|----|-------------------------------------------|
| 0  | Success.                                  |
| -1 | An error occurred ( <i>errno</i> is set). |

## Errors:

EPERM      The caller isn't `root`.

## Files:

`/etc/group`      The group database.

## Classification:

Unix

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## Caveats:

If *initgroups()* fails, it *doesn't* change the supplementary group access list.

The *getgrouplist()* function called by *initgroups()* is based on *getgrent()*. If the calling process uses *getgrent()*, the in-memory group structure is overwritten in the call to *initgroups()*.

## See also:

*getgroups()*, *getgrnam()*

**Synopsis:**

```
#include <stdlib.h>

char* initstate(unsigned int seed,
 char* state,
 size_t size);
```

**Arguments:**

*seed*     A starting point for the random-number sequence. This lets you restart the sequence at the same point.

*state*     The state array that you want to initialize.

*size*     The size, in bytes, of the state array; see below.

**Library:**

`libc`

**Description:**

The *initstate()* initializes the given state array for future use when generating pseudo-random numbers.

This function uses the *size* argument to determine what type of random-number generator to use; the larger the state array, the more random the numbers. Values for the amount of state information are 8, 32, 64, 128, and 256 bytes. Other values greater than 8 bytes are rounded down to the nearest one of these values. For values smaller than 8, *random()* uses a simple linear congruential random number generator.

Use this function in conjunction with the following:

*random()*     Generate a pseudo-random number using a default state.

*setstate()*   Specify the state of the pseudo-random number generator.

*srandom()* Set the seed used by the pseudo-random number generator.

If you haven't called *initstate()*, *random()* behaves as though you had called *initstate()* with a *seed* of 1 and a *size* of 128.

After initialization, you can restart a state array at a different point in one of these ways:

- Call *initstate()* with the desired seed, state array, and size of the array.
- Call *setstate()* with the desired state, then call *srandom()* with the desired seed. The advantage of using both of these functions is that the size of the state array doesn't have to be saved once it's initialized.

## Returns:

A pointer to the previous state array, or NULL if an error occurred.

## Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <time.h>

static char state1[32];

int main() {
 initstate(time(NULL), state1, sizeof(state1));
 setstate(state1);
 printf("%d0\n", random());
 return EXIT_SUCCESS;
}
```

## Classification:

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*drand48(), rand(), random(), setstate(), srand(), srandom()*

## ***input\_line()***

© 2004, QNX Software Systems Ltd.

*Get a string of characters from a file*

### **Synopsis:**

```
#include <stdio.h>

char* input_line(FILE* fp,
 char* buf,
 int bufsize);

extern int _input_line_max;
```

### **Arguments:**

|                |                                                                              |
|----------------|------------------------------------------------------------------------------|
| <i>fp</i>      | The file that you want to read from.                                         |
| <i>buf</i>     | A pointer to a buffer where the function can store the string that it reads. |
| <i>bufsize</i> | The size of the buffer, in bytes.                                            |

### **Library:**

libc

### **Description:**

The *input\_line()* function gets a string of characters from the file designated by *fp* and stores them in the array pointed to by *buf*. The *input\_line()* function stops reading characters when:

- end-of-file is reached
- a newline character is read
- *bufsize* - 1 characters have been read.

In addition, the *input\_line()* function buffers the last *\_input\_line\_max* lines internally. The *\_input\_line\_max* variable is defined in `<stdio.h>`. You can set it before calling *input\_line()* for the first time; its default value is 20. While the line is being read, the KEY\_UP and KEY\_DOWN keys can be used to move to the previous and next line respectively in a circular buffer of previously read lines. The newline character (`\n`) is replaced with the null character on input.



**Returns:**

A pointer to the input line. On end-of-file or on encountering an error reading from *fp*, NULL is returned and *errno* is set.

**Examples:**

```
#include <stdlib.h>
#include <stdio.h>

#define SIZ 256

int _input_line_max;

int main(void)
{
 FILE *fp;
 char *p,
 buf[SIZ];

 fp = stdin; /* Or any stream */
 _input_line_max = 25; /* set before 1st call */

 while((p = input_line(fp, buf, SIZ)) != NULL) {
 printf("%s\n", buf);
 fflush(stdout);
 }
 return EXIT_SUCCESS;
}
```

**Classification:**

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## ***InterruptAttach()*, *InterruptAttach\_r()*** © 2004, QNX Software Systems

Ltd.

---

*Attach an interrupt handler to an interrupt source*

### **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptAttach(int intr,
 const struct sigevent * (* handler)(void *, int),
 const void * area,
 int size,
 unsigned flags);

int InterruptAttach_r(int intr,
 const struct sigevent * (* handler)(void *, int),
 const void * area,
 int size,
 unsigned flags);
```

### **Arguments:**

|                |                                                                                                                                                            |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>intr</i>    | The interrupt that you want to attach a handler to; see “Interrupt vector numbers,” below.                                                                 |
| <i>handler</i> | A pointer to the handler function; see “Interrupt handler function,” below.                                                                                |
| <i>area</i>    | A pointer to a communications area in your process that the <i>handler</i> can assume is never paged out, or NULL if you don’t want a communications area. |
| <i>size</i>    | The size of the communications area.                                                                                                                       |
| <i>flags</i>   | Flags that specify how you want to attach the interrupt handler. For more information, see “Flags,” below.                                                 |

### **Library:**

`libc`

## Description:

The *InterruptAttach()* and *InterruptAttach\_r()* kernel calls attach the interrupt function *handler* to the hardware interrupt specified by *intr*. They automatically enable (i.e unmask) the interrupt level.

These functions are identical except in the way they indicate errors. See the Returns section for details.

Before calling either of these functions, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, the attachment fails with an error code of EPERM.

## Interrupt vector numbers

The interrupt values for *intr* are *logical interrupt vector numbers* grouped into related "interrupt classes" that generally correspond to a particular interrupt line on the CPU. The following interrupt classes are present on all QNX Neutrino systems:

`_NTO_INTR_CLASS_EXTERNAL`

Normal external interrupts (such as the ones generated by the **INTR** pin on x86 CPUs).

`_NTO_INTR_CLASS_SYNTHETIC`

Synthetic, kernel-generated interrupts.

`_NTO_INTR_SPARE` is usually the only `_NTO_INTR_CLASS_SYNTHETIC` interrupt you'll use; `_NTO_INTR_SPARE` is guaranteed not to match any valid logical interrupt vector number.

There can be additional interrupt classes defined for specific CPUs or embedded systems. For the interrupt assignments for specific boards, see the sample build files in

```
`${QNX_TARGET}/${PROCESSOR}/boot/build.
```

## Interrupts and startup code

The mapping of logical interrupt vector numbers is completely dependent on the implementor of the startup code.

Device drivers must:

- Let the user specify an interrupt number on the command line; don't use a hard-coded value. Eventually, the configuration manager will provide interrupt numbers for the device drivers.
- Store interrupt numbers in an **unsigned int** variable; don't assume an interrupt number fits into a byte.

## Typical x86 Interrupt vector numbers

The following list contains typical interrupt assignments for the 16 hardware interrupts on an x86-based PC using **startup-bios**:

| <b>Interrupt <i>intr</i></b> | <b>Description</b>                                              |
|------------------------------|-----------------------------------------------------------------|
| 0                            | A clock that runs at the resolution set by <i>ClockPeriod()</i> |
| 1                            | Keyboard                                                        |
| 2                            | Slave 8259 — you can't attach to this interrupt.                |
| 3                            | Com2                                                            |
| 4                            | Com1                                                            |
| 5                            | Net card / sound card / other                                   |
| 6                            | Floppy                                                          |
| 7                            | Parallel printer / sound card / other                           |
| 8                            |                                                                 |
| 9                            | Remapped interrupt 2                                            |
| 10                           |                                                                 |

*continued...*

| <b>Interrupt <i>intr</i></b> | <b>Description</b>        |
|------------------------------|---------------------------|
| 11                           |                           |
| 12                           |                           |
| 13                           | Co-processor              |
| 14                           | Primary disk controller   |
| 15                           | Secondary disk controller |




---

The interrupt assignments are different for other boards.

---

### **Interrupt handler function**

The function to call is specified by the *handler* argument. This function runs in the environment of your process. If a pager is running that swaps pages out of memory, It's possible for your *handler* to reference a variable in the process address space that isn't present. This results in a kernel shutdown.

The *area* and *size* arguments define a communications area in your process that the *handler* can assume is never paged out. This typically is a structure containing buffers and information needed by the *handler* and the process when it runs. In a paging system, lock the memory pointed to by *area* by calling *mlock()* before attaching the *handler*. In a nonpaging system, you can omit the call to *mlock()* (but you should still call it for compatibility with future versions of the OS).




---

The *area* argument can be NULL to indicate no communications area. If *area* is NULL, *size* should be 0.

---

The *handler* function's prototype is:

```
const struct sigevent* handler(void* area, int id);
```

Where *area* is a pointer to the *area* specified by the call to *InterruptAttach()*, and *id* is the ID returned by *InterruptAttach()*.

Follow the following guidelines when writing your handler:

- A temporary interrupt stack of limited depth is provided at interrupt time, so avoid placing large arrays or structures on the stack frame of the handler. It's safe to assume that about 200 bytes of stack are available.
- The interrupt handler runs asynchronously with the threads in the process. Any variables modified by the handler should be declared with the `volatile` keyword and modified with interrupts disabled or using the `atomic*()` functions in any thread and ISR.
- The interrupt handler should be kept as short as possible. If a significant amount of work needs to be done, the handler should deliver an event to awaken a thread to do the work.
- The handler can't call library routines that contain kernel calls *except* for `InterruptDisable()`, `InterruptEnable()`, `InterruptLock()`, `InterruptMask()`, `InterruptUnlock()`, and `InterruptUnmask()`.

The handler can call `TraceEvent()`, but not all modes are valid.

The return value of the *handler* function should be NULL or a pointer to a valid `sigevent` structure that the kernel delivers. These events are defined in `<signal.h>`.

Consider the following when choosing the event type:

- Message-driven processes that block in a receive loop using `MsgReceivev()` should consider using `SIGEV_PULSE` to trigger a pulse.
- Threads that block at a particular point in their code and don't go back to a common receive point should consider using `SIGEV_INTR` as the event notification type and `InterruptWait()` as the blocking call.




---

The thread that calls *InterruptWait()* *must* be the one that called *InterruptAttach()*.

---

- Using SIGEV\_SIGNAL, SIGEV\_SIGNAL\_CODE, SIGEV\_SIGNAL\_THREAD, or SIGEV\_THREAD is discouraged. It's less efficient than the other mechanisms for interrupt event delivery.

## Flags

The *flags* argument is a bitwise OR of the following values, or 0:

| Flag                    | Description                                                                                                       |
|-------------------------|-------------------------------------------------------------------------------------------------------------------|
| _NTO_INTR_FLAGS_END     | Put the new handler at the end of the list of existing handlers (for shared interrupts) instead of the start.     |
| _NTO_INTR_FLAGS_PROCESS | Associate the handler with the process instead of the attaching thread.                                           |
| _NTO_INTR_FLAGS_TRK_MSK | Track calls to <i>InterruptMask()</i> and <i>InterruptUnmask()</i> to make detaching the interrupt handler safer. |

## **\_NTO\_INTR\_FLAGS\_END**

The interrupt structure allows hardware interrupts to be shared. For example, if two processes take over the same physical interrupt, both handlers are invoked consecutively. When a handler attaches, it's placed in front of any existing handlers for that interrupt and is called first. You can change this behavior by setting the `_NTO_INTR_FLAGS_END` flag in the *flags* argument. This adds the handler at the end of any existing handlers. Although the Neutrino microkernel allows full interrupt sharing, your hardware might not.

For example, the ISA bus doesn't allow interrupt sharing, while the PCI bus does.

Processor interrupts are enabled during the execution of the handler. *Don't* attempt to talk to the interrupt controller chip. The operating system issues the end-of-interrupt command to the chip after processing all handlers at a given level.

The first process to attach to an interrupt unmask the interrupt. When the last process detaches from an interrupt, the system masks it.

If the thread that attached the interrupt handler terminates without detaching the handler, the kernel does it automatically.

#### **`_NTO_INTR_FLAGS_PROCESS`**

Adding `_NTO_INTR_FLAGS_PROCESS` to *flags* associates the interrupt handler with the *process* instead of the attaching thread. The interrupt handler is removed when the process exits, instead of when the attaching thread exits.

#### **`_NTO_INTR_FLAGS_TRK_MSK`**

The `_NTO_INTR_FLAGS_TRK_MSK` flag and the *id* argument to *InterruptMask()* and *InterruptUnmask()* let the kernel track the number of times a particular interrupt handler or event has been masked. Then, when an application detaches from the interrupt, the kernel can perform the proper number of unmask to ensure that the interrupt functions normally. This is important for shared interrupt levels.



You should always set `_NTO_INTR_FLAGS_TRK_MSK`.

---

### **Blocking states**

This call doesn't block.



## Returns:

The only difference between these functions is the way they indicate errors:

### *InterruptAttach()*

An interrupt function ID. If an error occurs, -1 is returned and *errno* is set.

### *InterruptAttach\_r()*

An interrupt function ID. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

Use the function ID with the *InterruptDetach()* function to detach this interrupt handler.

## Errors:

|        |                                                                        |
|--------|------------------------------------------------------------------------|
| EAGAIN | All kernel interrupt entries are in use.                               |
| EFAULT | A fault occurred when the kernel tried to access the buffers provided. |
| EINVAL | The value of <i>intr</i> isn't a valid interrupt number.               |
| EPERM  | The process doesn't have I/O privileges.                               |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

If you're writing a resource manager and using the *resmgr\_\**() functions with multiple threads, a thread that attaches to an interrupt *must* use `_NTO_INTR_FLAGS_PROCESS` in the *flags* argument when calling *InterruptAttach()*.

If your interrupt handler isn't SMP-safe, you must lock it to one processor using:

```
ThreadCtl(_NTO_TCTL_RUNMASK, ...);
```

## See also:

*atomic\_add()*, *atomic\_clr()*, *atomic\_set()*, *atomic\_sub()*,  
*atomic\_toggle()*, *InterruptAttachEvent()*, *InterruptDetach()*,  
*InterruptDisable()*, *InterruptEnable()*, *InterruptLock()*,  
*InterruptMask()*, *InterruptUnlock()*, *InterruptUnmask()*,  
*InterruptWait()*, *mlock()*, **sigevent**, *ThreadCtl()*, *TraceEvent()*

## ***InterruptAttachEvent()*, *InterruptAttachEvent\_r()***

*Attach an event to an interrupt source*

### **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptAttachEvent(
 int intr,
 const struct sigevent* event,
 unsigned flags);

int InterruptAttachEvent_r(
 int intr,
 const struct sigevent* event,
 unsigned flags);
```

### **Arguments:**

*intr*      The *interrupt vector number* that you want to attach an event to; for more information, see “Interrupt vector numbers” in the documentation for *InterruptAttach()*.

*event*     A pointer to the **sigevent** structure that you want to be delivered when this interrupt occurs.

*flags*     Flags that specify how you want to attach the interrupt handler. For more information, see “Flags,” below.

### **Library:**

`libc`

### **Description:**

The *InterruptAttachEvent()* and *InterruptAttachEvent\_r()* kernel calls attach the given event to the hardware interrupt specified by *intr*. They automatically enable (i.e unmask) the interrupt level.

The *InterruptAttachEvent()* and *InterruptAttachEvent\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

## *InterruptAttachEvent()*, *InterruptAttachEvent\_r()* ©

2004, QNX Software Systems Ltd.

---

Before calling either of these functions, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV when it calls *InterruptUnlock()*.

To prevent infinite interrupt recursion, the kernel automatically does an *InterruptMask()* for *intr* when delivering the event. After the interrupt-handling thread has dealt with the event, it must call *InterruptUnmask()* to reenale the interrupt.

Consider the following when choosing an event type:

- Message-driven processes that block in a receive loop using *MsgReceivev()* should consider using SIGEV\_PULSE to trigger a channel.
- Threads that block at a particular point in their code and don't go back to a common receive point, should consider using SIGEV\_INTR as the event notification type and *InterruptWait()* as the blocking call.



---

The thread that calls *InterruptWait()* *must* be the one that called *InterruptAttachEvent()*.

---

- Using SIGEV\_SIGNAL, SIGEV\_SIGNAL\_CODE, or SIGEV\_SIGNAL\_THREAD is discouraged. It is less efficient than the other mechanisms for interrupt event delivery.

### Flags

The *flags* argument is a bitwise OR of the following values, or 0:

## *InterruptAttachEvent()*, *InterruptAttachEvent\_r()*

| Flag                                 | Description                                                                                                       |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <code>_NTO_INTR_FLAGS_END</code>     | Put the new event at the end of the list of existing events instead of the start.                                 |
| <code>_NTO_INTR_FLAGS_PROCESS</code> | Associate the event with the process instead of the attaching thread.                                             |
| <code>_NTO_INTR_FLAGS_TRK_MSK</code> | Track calls to <i>InterruptMask()</i> and <i>InterruptUnmask()</i> to make detaching the interrupt handler safer. |

### **`_NTO_INTR_FLAGS_END`**

The interrupt structure allows hardware interrupts to be shared. For example if two processes call *InterruptAttachEvent()* for the same physical interrupt, both events are sent consecutively. When an event attaches, it's placed in front of any existing events for that interrupt and is delivered first. You can change this behavior by setting the `_NTO_INTR_FLAGS_END` flag in the *flags* argument. This adds the event at the end of any existing events.

### **`_NTO_INTR_FLAGS_PROCESS`**

Adding `_NTO_INTR_FLAGS_PROCESS` to *flags* associates the interrupt event with the *process* instead of the attaching thread. The interrupt event is removed when the process exits, instead of when the attaching thread exits.



The kernel automatically attempts to set the `_NTO_INTR_FLAGS_PROCESS` flag if the event is directed at the process in general (for `SIGEV_SIGNAL`, `SIGEV_SIGNAL_CODE`, and `SIGEV_PULSE` events).

## *InterruptAttachEvent()*, *InterruptAttachEvent\_r()* ©

2004, QNX Software Systems Ltd.

---

### **.\_NTO\_INTR\_FLAGS\_TRK\_MSK**

The `._NTO_INTR_FLAGS_TRK_MSK` flag and the *id* argument to *InterruptMask()* and *InterruptUnmask()* let the kernel track the number of times a particular interrupt handler or event has been masked. Then, when an application detaches from the interrupt, the kernel can perform the proper number of unmask to ensure that the interrupt functions normally. This is important for shared interrupt levels.



---

You should always set `._NTO_INTR_FLAGS_TRK_MSK`.

---

### **Advantages & disadvantages**

*InterruptAttachEvent()* has several advantages over *InterruptAttach()*:

- Less work is done at interrupt time (you avoid the context switch necessary to map in an interrupt handler).
- Interrupt handling code runs at the thread's priority, which lets you specify the priority of the interrupt handling.
- You can use process-level debugging on your interrupt handler code.

There are also some disadvantages:

- There might be a delay before the interrupt handling code runs (until the thread is scheduled to run).
- For multiple devices sharing an event, the amount of time spent with the interrupt masked increases.

You can freely mix calls to *InterruptAttach()* and *InterruptAttachEvent()* for a particular interrupt.

### **Blocking states**

This call doesn't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *InterruptAttachEvent()*

An interrupt function ID. If an error occurs, -1 is returned and *errno* is set.

### *InterruptAttachEvent\_r()*

An interrupt function ID. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

Use the ID with *InterruptDetach()* to detach this interrupt event.

## Errors:

|        |                                                                        |
|--------|------------------------------------------------------------------------|
| EAGAIN | All kernel interrupt entries are in use.                               |
| EFAULT | A fault occurred when the kernel tried to access the buffers provided. |
| EINVAL | The value of <i>intr</i> isn't a valid interrupt number.               |
| EPERM  | The process doesn't have superuser capabilities.                       |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## *InterruptAttachEvent()*, *InterruptAttachEvent\_r()* ©

2004, QNX Software Systems Ltd.

---

### **See also:**

*InterruptAttach()*, *InterruptDetach()*, *InterruptLock()*,  
*InterruptMask()*, *InterruptUnlock()*, *InterruptUnmask()*,  
*InterruptWait()*, **sigevent**



## ***InterruptDetach(), InterruptDetach\_r()***

*Detach an interrupt handler by ID*

### **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptDetach(int id);

int InterruptDetach_r(int id);
```

### **Arguments:**

*id* The value returned by *InterruptAttach()*, *InterruptAttachEvent()*, or *InterruptHookIdle()*.

### **Library:**

libc

### **Description:**

These kernel calls detach the interrupt handler specified by the *id* argument. If, after detaching, no thread is attached to the interrupt then the interrupt is masked off.

The *InterruptDetach()* and *InterruptDetach\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

Before calling either of these functions, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV when it calls *InterruptUnlock()*.

### **Blocking states**

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *InterruptDetach()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

### *InterruptDetach\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

- EINVAL     The value of *id* doesn't exist for this process.
- EPERM     The process doesn't have superuser capabilities.

## Classification:

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*InterruptAttach()*, *InterruptAttachEvent()*, *InterruptHookIdle()*,  
*InterruptUnlock()*

### **Synopsis:**

```
#include <sys/neutrino.h>

void InterruptDisable(void);
```

### **Library:**

libc

### **Description:**

The *InterruptDisable()* function disables all hardware interrupts. You can call it from a thread or from an interrupt handler. Before calling this function, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```



---

Any kernel call results in the re-enabling of interrupts, and many library routines are built on kernel calls. Masked interrupts are not affected.

---

If the thread doesn't do this, it might SIGSEGV when *InterruptUnlock()* is called.

Reenable the interrupts by calling *InterruptEnable()*.



---

**CAUTION:** Since this function disables all hardware interrupts, take care to reenable them as quickly as possible. Failure to do so may result in increased interrupt latency and nonrealtime performance.

---

Use *InterruptDisable()* instead of an inline `cli` to ensure hardware portability with non-x86 CPUs.



---

Use *InterruptLock()* and *InterruptUnlock()* instead of *InterruptDisable()* and *InterruptEnable()*. The *InterruptLock()* and *InterruptUnlock()* functions perform the intended function on SMP hardware, and allow your interrupt thread to run on any processor in the system.

---

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*InterruptEnable()*, *InterruptLock()*, *InterruptMask()*,  
*InterruptUnlock()*, *InterruptUnmask()*, *ThreadCtl()*

## **Synopsis:**

```
#include <sys/neutrino.h>

void InterruptEnable(void);
```

## **Library:**

libc

## **Description:**

The *InterruptEnable()* function enables all hardware interrupts. You can call it from a thread or from an interrupt handler. Before calling this function, the thread must request I/O priority by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV when *InterruptUnlock()* is called.

You should call this function as quickly as possible after calling *InterruptDisable()*.



---

Use *InterruptLock()* and *InterruptUnlock()* instead of *InterruptDisable()* and *InterruptEnable()*. The *InterruptLock()* and *InterruptUnlock()* functions perform the intended function on SMP hardware, and allow your interrupt thread to run on any processor in the system.

---

## **Classification:**

QNX Neutrino

### **Safety**

---

Cancellation point No

*continued...*

## **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

## **See also:**

*InterruptDisable(), InterruptLock(), InterruptMask(),  
InterruptUnlock(), InterruptUnmask(), ThreadCtl()*

## Synopsis:

```
#include <sys/neutrino.h>

int InterruptHookIdle(
 void (*handler) (uint64_t *, struct qtime_entry *),
 unsigned flags);
```

## Arguments:

*handler*     A pointer to the handler function; see below.

*flags*        Flags that specify how you want to attach the interrupt handler. For more information, see “Flags,” below.

## Library:

`libc`

## Description:

The *InterruptHookIdle()* kernel call attaches the specified interrupt *handler* to the “idle” interrupt, which is called when the system is idle. This is typically used to implement power management features.

The arguments to the *handler* functions are:

`uint64_t*`     A pointer to the time, in nanoseconds, when the next timer will expire.

`struct qtime_entry *`  
A pointer to the section of the system page with the time information, including the current time of day.

The simplest idle handler consists of a `halt` instruction.

## Flags

The *flags* argument is a bitwise OR of the following values, or 0:

| Flag                                 | Description                                                                                                       |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <code>_NTO_INTR_FLAGS_END</code>     | Put the new handler at the end of the list of existing handlers (for shared interrupts) instead of the start.     |
| <code>_NTO_INTR_FLAGS_PROCESS</code> | Associate the handler with the process instead of the attaching thread.                                           |
| <code>_NTO_INTR_FLAGS_TRK_MSK</code> | Track calls to <i>InterruptMask()</i> and <i>InterruptUnmask()</i> to make detaching the interrupt handler safer. |

### `_NTO_INTR_FLAGS_END`

The interrupt structure allows hardware interrupts to be shared. For example, if two processes take over the same physical interrupt, both handlers are invoked consecutively. When a handler attaches, it's placed in front of any existing handlers for that interrupt and is called first. You can change this behavior by setting the `_NTO_INTR_FLAGS_END` flag in the *flags* argument. This adds the handler at the end of any existing handlers.

Processor interrupts are enabled during the execution of the handler. *Don't* attempt to talk to the interrupt controller chip. The end of interrupt command is issued to the chip by the operating system after processing all handlers at a given level.

The first process to attach to an interrupt unmask the interrupt. When the last process detaches from an interrupt, the system masks it.

If the thread that attached the interrupt handler terminates without detaching the handler, the kernel does it automatically.



### **.\_NTO\_INTR\_FLAGS\_PROCESS**

Adding `._NTO_INTR_FLAGS_PROCESS` to *flags* associates the interrupt handler with the *process* instead of the attaching thread. The interrupt handler is removed when the process exits, instead of when the attaching thread exits.

### **.\_NTO\_INTR\_FLAGS\_TRK\_MSK**

The `._NTO_INTR_FLAGS_TRK_MSK` flag and the *id* argument to *InterruptMask()* and *InterruptUnmask()* let the kernel track the number of times a particular interrupt handler or event has been masked. Then, when an application detaches from the interrupt, the kernel can perform the proper number of unmask to ensure that the interrupt functions normally. This is important for shared interrupt values.

### **Blocking states**

This call doesn't block.

### **Returns:**

An interrupt function ID, or -1 if an error occurs (*errno* is set).

Use the returned value with the *InterruptDetach()* function to detach this interrupt handler.

### **Errors:**

EAGAIN All kernel interrupt entries are in use.

EPERM The process doesn't have superuser capabilities.

### **Classification:**

QNX Neutrino

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*InterruptAttach(), InterruptAttachEvent(), InterruptDetach(),  
InterruptHookTrace()*

*Attach the pseudo interrupt handler that the instrumented module uses*

### **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptHookTrace(
 const struct sigevent * (* handler)(int),
 unsigned flags);
```

### **Arguments:**

*handler*     A pointer to the handler function.

*flags*        Flags that specify how you want to attach the interrupt handler.

### **Library:**

`libc`

### **Description:**

The *InterruptHookTrace()* kernel call attaches the pseudo interrupt handler *handle* that the instrumented module uses.



---

This function requires the instrumented kernel. For more information, see the documentation for the System Analysis Toolkit (SAT).

---

### **Returns:**

An interrupt function ID, or -1 if an error occurs (*errno* is set).

### **Errors:**

EAGAIN        All kernel interrupt entries are in use.

EFAULT        A fault occurred when the kernel tried to access the buffers provided.

EPERM         The process doesn't have superuser capabilities.

ENOTSUP       The kernel is not instrumented.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*InterruptAttach(), TraceEvent()*

### **Synopsis:**

```
#include <sys/neutrino.h>

void InterruptLock(intrspin_t* spinlock);
```

### **Arguments:**

*spinlock*     The spinlock (a variable shared between the interrupt handler and a thread) to use.



---

If *spinlock* isn't a **static** variable, you must initialize it by calling:

```
memset(spinlock, 0, sizeof(*spinlock));
```

before using it with *InterruptLock()*.

---

### **Library:**

**libc**

### **Description:**

The *InterruptLock()* function guards a critical section by locking the specified *spinlock*. You can call this function from a thread or from an interrupt handler. Before calling this function, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV when *InterruptUnlock()* is called.

This function tries to acquire the *spinlock* (a variable shared between the interrupt handler and a thread) while interrupts are disabled. The code spins in a tight loop until the lock is acquired. It's important to release the lock as soon as possible. Typically, this is a few lines of code without any loops:

```
InterruptLock(&spinner);

/* ... critical section */

InterruptUnlock(&spinner);
```

*InterruptLock()* solves a common need in many realtime systems to protect access to shared data structures between an interrupt handler and the thread that owns the handler. The traditional POSIX primitives used between threads aren't available for use by an interrupt handler.

The *InterruptLock()* and *InterruptUnlock()* functions work on single-processor or multiprocessor machines.



---

Any kernel call results in the re-enabling of interrupts, and many library routines are built on kernel calls. Masked interrupts are not affected.

---

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*InterruptDisable()*, *InterruptEnable()*, *InterruptMask()*,  
*InterruptUnlock()*, *InterruptUnmask()*, *ThreadCtl()*

## Synopsis:

```
#include <sys/neutrino.h>

int InterruptMask(int intr,
 int id);
```

## Arguments:

*intr*     The interrupt you want to mask.

*id*       The value returned by *InterruptAttach()* or *InterruptAttachEvent()*, or -1 if you don't want the kernel to track interrupt maskings and unmaskings for each handler.



---

The *id* is ignored unless you use the `_NTO_INTR_FLAGS_TRK_MSK` flag when you attach the handler.

---

## Library:

`libc`

## Description:

The *InterruptMask()* kernel call disables the hardware interrupt specified by *intr* for the handler specified by *id*. You can call this function from a thread or from an interrupt handler. Before calling this function, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV when *InterruptUnmask()* is called.

Reenable the interrupt by calling *InterruptUnmask()*.

The kernel automatically enables an interrupt when the first handler attaches to it using *InterruptAttach()* and disables it when the last handler detaches.

This call is often used when a device presents a level-sensitive interrupt to the system that can't be easily cleared in the interrupt handler. Since the interrupt is level-sensitive, you can't exit the handler with the interrupt line active and unmasked. *InterruptMask()* lets you mask the interrupt in the handler and schedule a thread to do the real work of communicating with the device to clear the source. Once cleared, the thread should call *InterruptUnmask()* to reenable this interrupt.

To disable all hardware interrupts, use the *InterruptLock()* function.



---

To ensure hardware portability, use *InterruptMask()* instead of writing code that talks directly to the interrupt controller.

---

Calls to *InterruptMask()* are nested; the interrupt isn't unmasked until *InterruptUnmask()* has been called once for every call to *InterruptMask()*.

## Returns:

The current mask level count for success; or -1 if an error occurs (*errno* is set).

## Errors:

EINVAL     The value of *intr* isn't a supported hardware interrupt.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*InterruptAttach(), InterruptDisable(), InterruptEnable(),  
InterruptLock(), InterruptUnlock(), InterruptUnmask(), ThreadCtl()*

# ***InterruptUnlock()***

© 2004, QNX Software Systems Ltd.

*Release a critical section in an interrupt handler*

## **Synopsis:**

```
#include <sys/neutrino.h>

void InterruptUnlock(intrspin_t* spinlock);
```

## **Arguments:**

*spinlock* The spinlock (a variable shared between the interrupt handler and a thread) used in a call to *InterruptLock()* to lock the handler.

## **Library:**

libc

## **Description:**

The *InterruptUnlock()* function releases a critical section by unlocking the specified *spinlock*, reenabling interrupts. You can call this function from a thread or from an interrupt handler.

Before calling this function, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*InterruptDisable(), InterruptEnable(), InterruptLock(),  
InterruptMask(), InterruptUnmask(), ThreadCtl()*

# ***InterruptUnmask()***

© 2004, QNX Software Systems Ltd.

*Enable a hardware interrupt*

## **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptUnmask(int intr,
 int id);
```

## **Arguments:**

*intr*     The interrupt you want to unmask.

*id*       The value returned by *InterruptAttach()* or *InterruptAttachEvent()*, or -1 if you don't want the kernel to track interrupt maskings and unmaskings for each handler.



---

The *id* is ignored unless you use the `_NTO_INTR_FLAGS_TRK_MSK` flag when you attach the handler.

---

## **Library:**

`libc`

## **Description:**

The *InterruptUnmask()* kernel call enables the hardware interrupt specified by *intr* for the interrupt handler specified by *intr* for the handler specified by *id* when the mask count reaches zero. You can call this function from a thread or from an interrupt handler. Before calling this function, the thread must request I/O privity by calling:

```
ThreadCtl(_NTO_TCTL_IO, 0);
```

If the thread doesn't do this, it might SIGSEGV.

Calls to *InterruptMask()* are nested; the interrupt isn't unmasked until *InterruptUnmask()* has been called once for every call to *InterruptMask()*.

## **Returns:**

The current mask count, or -1 if an error occurs (*errno* is set).

## **Errors:**

EINVAL      Not a supported hardware interrupt *intr*.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*InterruptAttach()*, *InterruptDisable()*, *InterruptEnable()*,  
*InterruptLock()*, *InterruptMask()*, *InterruptUnlock()* *ThreadCtl()*

# ***InterruptWait(), InterruptWait\_r()***

© 2004, QNX Software Systems Ltd.

*Wait for a hardware interrupt*

## **Synopsis:**

```
#include <sys/neutrino.h>

int InterruptWait(int flags,
 const uint64_t * timeout);

int InterruptWait_r(int flags,
 const uint64_t * timeout);
```

## **Arguments:**

*flags*        This should currently be 0.

*timeout*      This should currently be NULL. This may change in future versions.



---

Use *TimerTimeout()* to achieve a timeout.

---

## **Library:**

libc

## **Description:**

These kernel calls wait for a hardware interrupt. The calling thread should have attached a handler to the interrupt, by calling *InterruptAttach()* or *InterruptAttachEvent()*. The call to *InterruptWait()* or *InterruptWait\_r()* blocks waiting for an interrupt handler to return an event with notification type SIGEV\_INTR (i.e. a hardware interrupt).

The *InterruptWait()* and *InterruptWait\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

If the notification event occurs before *InterruptWait()* is called, a pending flag is set. When *InterruptWait()* is called, the flag is checked; if set, it's cleared and the call immediately returns with success.

**Blocking states**

STATE\_INTR      The thread is waiting for an interrupt handler to return a SIGEV\_INTR event.

**Returns:**

The only difference between these functions is the way they indicate errors:

*InterruptWait()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*InterruptWait\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

**Errors:**

EINTR            The call was interrupted by a signal.

ENOTSUP        The reserved arguments aren't NULL.

ETIMEDOUT      A kernel timeout unblocked the call. See *TimerTimeout()*.

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*InterruptAttach(), InterruptAttachEvent(), TimerTimeout()*



**Synopsis:**

```
#include <x86/v86.h>

int _intr_v86(int swi,
 struct _v86reg* regs,
 void* data,
 int datasize);
```

**Arguments:**

*swi*            The software interrupt that you want to execute.

*regs*           A pointer to a `_v86reg` structure that specifies the values you want to use for the registers on entry to real mode; see below.

*data*           A pointer to the data that you want to copy into memory; see below.

*datasize*      The size of the data, in bytes.

**Library:**

`libc`

**Description:**

The `_intr_v86()` function executes the real-mode software interrupt specified by *swi* in virtual 8086 mode. This allows access to the ROM BIOS functions that are designed to run in 16-bit real mode. Two common examples are:

| <b>Interrupt</b>     | <b>Description</b> |
|----------------------|--------------------|
| <code>int 10h</code> | Video BIOS         |
| <code>int 1ah</code> | PCI                |

BIOS calls (such as `int 13h`, disk I/O) that require hardware interrupts to be directed at their code aren't supported.

Upon entry to real mode, the registers are loaded from *regs*. The segment registers and any pointers should address a 2K communication area located at offset `0:800h` in real memory. The buffer *data* of length *datasize* is copied to this area just before real mode is entered and copied back when the call completes. At this point *regs* is also updated to contain the values of the real-mode registers.

You should set the **DS**, **ES**, **FS** and **GS** segment registers to 0. The values in the **CS:IP**, and **SS:SP** registers are ignored and are set by the kernel. The stack provided is about 500 bytes in size.

The layout of real mode memory is described by the structure `_v86_memory` in `<x86/v86.h>`.

When a thread enters virtual 8086 mode, all threads in the system continue to be scheduled based upon their priority, including the calling thread. While in virtual 8086 mode, full access to IO ports and interrupt enable and disable are allowed. Only one thread may enter virtual 8086 mode at a time.

This function fails if the calling process doesn't have an effective user ID of `root` (*euid* 0).

## Returns:

- 0 Success.
- 1 An error occurred; *errno* is set.

## Errors:

- EPERM The calling thread didn't have an effective user ID of `root`.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <stddef.h>
#include <errno.h>
#include <x86/v86.h>

struct _v86reg reg;

int main(void) {
 char buf[4];

 /* Equipment call */
 printf("int 12\n");
 memset(®, 0, sizeof(reg));
 _intr_v86(0x12, ®, NULL, 0);
 printreg();
 sleep(5);

 /* Enter 40 column text mode */
 printf("int 10 ah=00h al=00h\n");
 memset(®, 0, sizeof(reg));
 _intr_v86(0x10, ®, NULL, 0);
 printreg();
 sleep(5);

 /* Enter 80 column text mode */
 printf("int 10 ah=00h al=02h\n");
 memset(®, 0, sizeof(reg));
 reg.eax = 2;
 _intr_v86(0x10, ®, NULL, 0);
 printreg();
 sleep(5);

 /* Write a string from memory */
 printf("int 10 ah=13h al=00h\n");
 strcpy(buf, "Hi!");
 memset(®, 0, sizeof(reg));
 reg.eax = 0x1300;
 reg.es = 0;
 reg.ebp = offsetof(struct _v86_memory, userdata);
 reg.ecx = strlen(buf);
 reg.edx = 0;
 reg.ebx = 0x0007;
 _intr_v86(0x10, ®, buf, strlen(buf));
 printreg();
 sleep(5);

 return EXIT_SUCCESS;
}
```

```
printreg() {

 printf("eax=%-8x ebx=%-8x ecx=%-8x edx=%-8x\n",
 reg.eax, reg.ebx, reg.ecx, reg.edx);
 printf("esi=%-8x edi=%-8x ebp=%-8x esp=%-8x\n",
 reg.esi, reg.edi, reg.ebp, reg.esp);
 printf(" ds=%-8x es=%-8x fs=%-8x gs=%-8x\n",
 reg.ds, reg.es, reg.fs, reg.gs);
 printf("efl=%-8x\n\n",
 reg.efl);
}
```

## **Classification:**

QNX Neutrino (x86 only)

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Synopsis:**

```

struct _io_connect {
 uint16_t type;
 uint16_t subtype;
 uint32_t file_type;
 uint16_t reply_max;
 uint16_t entry_max;
 uint32_t key;
 uint32_t handle;
 uint32_t ioflag;
 uint32_t mode;
 uint16_t sflag;
 uint16_t access;
 uint16_t zero;
 uint16_t path_len;
 uint8_t eflag;
 uint8_t extra_type;
 uint16_t extra_len;
 char path[1];
};

```

**Description:**

The `_io_connect` structure is used to describe a connect message that a resource manager receives and sends.

The members include:

|                |                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>    | IO_CONNECT                                                                                                                                                                                                                                                                                                                                                       |
| <i>subtype</i> | The type of connection that the message concerns; one of: <ul style="list-style-type: none"> <li>• <code>_IO_CONNECT_COMBINE</code> — combine with an I/O message.</li> <li>• <code>_IO_CONNECT_COMBINE_CLOSE</code> — combine with I/O message and always close.</li> <li>• <code>_IO_CONNECT_OPEN</code></li> <li>• <code>_IO_CONNECT_UNLINK</code></li> </ul> |

- `_IO_CONNECT_RENAME`
- `_IO_CONNECT_MKNOD`
- `_IO_CONNECT_READLINK`
- `_IO_CONNECT_LINK`
- `_IO_CONNECT_RSVD_UNBLOCK` — place holder in the jump table.
- `_IO_CONNECT_MOUNT`

*file\_type*

The file type; one of the following (defined in `<sys/ftype.h>`):

- `_FTYPE_ANY` — the path name can be anything.
- `_FTYPE_LINK` — reserved for the Process Manager.
- `_FTYPE_MOUNT` — receive mount requests on the path (*path* must be NULL).
- `_FTYPE_QUEUE` — reserved for a mqueue manager.
- `_FTYPE_PIPE` — reserved for a pipe manager.
- `_FTYPE_SEM` — reserved for a semaphore manager.
  
- `_FTYPE_SHMEM` — reserved for a shared memory object.
- `_FTYPE_SOCKET` — reserved for a socket manager.
- `_FTYPE_SYMLINK` — reserved for the Process Manager.

*reply\_max*

The maximum length of the reply message.

*entry\_max*

The maximum number of `_io_connect_entry` structures that the resource manager is willing to accept. If a path could reference more than one resource manager, it returns a list of `_io_connect_entry` structures referring to the overlapping resource managers.

*key*

Reserved.

*handle* The handle returned by *resmgr\_attach()*.

*ioflag* One of:

- O\_RDONLY — open for reading only.
- O\_RDWR — open for reading and writing. Opening a FIFO for read-write is unsupported.
- O\_WRONLY — open for writing only.

You can also specify any combination of the remaining flags in the value of *ioflag*:

- O\_APPEND — if set, the file offset is set to the end of the file prior to each write.
- O\_CLOEXEC — close the file descriptor on execution.
- O\_CREAT — create the file.
- O\_DSYNC — if set, this flag affects subsequent I/O calls; each call to *write()* waits until all data is successfully transferred to the storage device such that it's readable on any subsequent open of the file (even one that follows a system failure) in the absence of a failure of the physical storage medium. If the physical storage medium implements a non-write-through cache, then a system failure may be interpreted as a failure of the physical storage medium, and data may not be readable even if this flag is set and the *write()* indicates that it succeeded.
- O\_EXCL — if you set both O\_EXCL and O\_CREAT, *open()* fails if the file exists. The check for the existence of the file and the creation of the file if it doesn't exist are atomic; no other process that's attempting the same operation with the same filename at the same time will succeed. Specifying O\_EXCL without O\_CREAT has no effect.
- O\_LARGEFILE — allow the file offset to be 64 bits long.

- O\_NOCTTY — if set, and *path* identifies a terminal device, the *open()* function doesn't cause the terminal device to become the controlling terminal for the process.
- O\_NONBLOCK — don't block.
- O\_REALIDS — use the real **uid/gid** for permissions checking.
- O\_RSYPNC — read I/O operations on the file descriptor complete at the same level of integrity as specified by the O\_DSYPNC and O\_SYPNC flags.
- O\_SYNC — if set, this flag affects subsequent I/O calls; each call to *read()* or *write()* is complete only when both the data has been successfully transferred (either read or written) and all file system information relevant to that I/O operation (including that required to retrieve said data) is successfully transferred, including file update and/or access times, and so on. See the discussion of a successful data transfer in O\_DSYPNC, above.
- O\_TRUNC — if the file exists and is a regular file, and the file is successfully opened O\_WRONLY or O\_RDWR, the file length is truncated to zero and the mode and owner are left unchanged. O\_TRUNC has no effect on FIFO or block or character special files or directories. Using O\_TRUNC with O\_RDONLY has no effect.

*mode*

One of:

- S\_IFBLK — block special.
- S\_IFCHR — character special.
- S\_IFDIR — directory.
- S\_IFIFO — FIFO special.
- S\_IFLNK — symbolic link.
- S\_IFMT — type of file.



- S\_IFNAM — special named file.
- S\_IFREG — regular.
- S\_IFSOCK — socket.

*sflag*

How the client wants the file to be shared; a combination of the following bits:

- SH\_COMPAT — set compatibility mode.
- SH\_DENYRW — prevent read or write access to the file.
- SH\_DENYWR — prevent write access to the file.
- SH\_DENYRD — prevent read access to the file.
- SH\_DENYNO — permit both read and write access to the file.

*access*

The access permissions for the file or directory, specified as a combination of the following bits (defined in `<sys/stat.h>`):

| <b>Owner</b> | <b>Group</b> | <b>Others</b> | <b>Permission</b>                                                                                                                   |
|--------------|--------------|---------------|-------------------------------------------------------------------------------------------------------------------------------------|
| S_IRUSR      | S_IRGRP      | S_IROTH       | Read                                                                                                                                |
| S_IRWXU      | S_IRWXG      | S_IRWXO       | Read, write, execute/search. A bitwise inclusive OR of the other three constants.<br>(S_IRWXU is OR of IRUSR, S_IWSUR and S_IXUSR.) |
| S_IWUSR      | S_IWGRP      | S_IWOTH       | Write                                                                                                                               |
| S_IXUSR      | S_IXGRP      | S_IXOTH       | Execute/search                                                                                                                      |

The following bits define miscellaneous permissions used by other implementations:

|                   | <b>Bit</b> | <b>Equivalent</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                   | S_IEXEC    | S_IXUSR                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                   | S_IREAD    | S_IRUSR                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                   | S_IWRITE   | S_IWUSR                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <i>path_len</i>   |            | The length of the <i>path</i> member.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>eflag</i>      |            | Extended flags: <ul style="list-style-type: none"><li>• <code>_IO_CONNECT_EFLAG_DIR</code> — the path referenced a directory.</li><li>• <code>_IO_CONNECT_EFLAG_DOT</code> — the last component of a path was <code>.</code> or <code>..</code> (i.e. the current or parent directory).</li></ul>                                                                                                                                                                                                                                                                                                                                                     |
| <i>extra_type</i> |            | One of: <ul style="list-style-type: none"><li>• <code>_IO_CONNECT_EXTRA_NONE</code></li><li>• <code>_IO_CONNECT_EXTRA_LINK</code></li><li>• <code>_IO_CONNECT_EXTRA_SYMLINK</code></li><li>• <code>_IO_CONNECT_EXTRA_QUEUE</code></li><li>• <code>_IO_CONNECT_EXTRA_PHOTON</code></li><li>• <code>_IO_CONNECT_EXTRA_SOCKET</code></li><li>• <code>_IO_CONNECT_EXTRA_SEM</code></li><li>• <code>_IO_CONNECT_EXTRA_RESMGR_LINK</code></li><li>• <code>_IO_CONNECT_EXTRA_PROC_SYMLINK</code></li><li>• <code>_IO_CONNECT_EXTRA_RENAME</code></li><li>• <code>_IO_CONNECT_EXTRA_MOUNT</code></li><li>• <code>_IO_CONNECT_EXTRA_MOUNT_OCB</code></li></ul> |
| <i>extra_len</i>  |            | The length of any extra data included in the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>path</i>       |            | The path that the client is trying to connect to, relative to the resource manager's mountpoint.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## Classification:

QNX Neutrino

## See also:

`_io_connect_fstype_reply`, `_io_connect_link_reply`,  
`resmgr_connect_funcs_t`

“The `_IO_OPEN` message for filesystems” in the Writing a Resource Manager chapter of the *Programmer's Guide*.

## **\_io\_connect\_ftype\_reply**

© 2004, QNX Software Systems Ltd.

*Structure of a connect message giving a status and a file type*

### **Synopsis:**

```
struct _io_connect_ftype_reply {
 uint16_t
 uint16_t
 uint32_t
};
```

**status;**  
**reserved;**  
**file\_type;**

### **Description:**

A resource manager uses the `_io_connect_ftype_reply` structure to send a status and a file type to a client that has sent a connect message.

The members include:

*status* Typically one of the *errno* values.

*file\_type* The file type; one of the following (defined in `<sys/ftype.h>`):

- `_FTYPE_ANY` — the path name can be anything.
- `_FTYPE_LINK` — reserved for the Process Manager.
- `_FTYPE_MOUNT` — receive mount requests on the path (*path* must be NULL).
- `_FTYPE_MQUEUE` — reserved for a mqueue manager.
- `_FTYPE_PIPE` — reserved for a pipe manager.
- `_FTYPE_SEM` — reserved for a semaphore manager.
- `_FTYPE_SHMEM` — reserved for a shared memory object.
- `_FTYPE_SOCKET` — reserved for a socket manager.
- `_FTYPE_SYMLINK` — reserved for the Process Manager.

**Classification:**

QNX Neutrino

**See also:**

`_io_connect`, `_io_connect_link_reply`,  
`resmgr_connect_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## **\_io\_connect\_link\_reply**

© 2004, QNX Software Systems Ltd.

*Structure of a connect message that redirects a client to another resource*

### **Synopsis:**

```
struct _io_connect_link_reply {
 uint32_t reserved1;
 uint32_t file_type;
 uint8_t eflag;
 uint8_t reserved2[1];
 uint16_t chroot_len;
 uint32_t umask;
 uint16_t nentries;
 uint16_t path_len;
 /*
 struct _io_connect_entry server[nentries];
 char path[path_len];
 or
 struct _server_info info;
 io_?_t msg;
 */
};
```

### **Description:**

A resource manager uses the `_io_connect_link_reply` structure in a reply to a client that redirects the client to another resource. The members include:

- file\_type*      The file type; one of the following (defined in `<sys/ftype.h>`):
- `_FTYPE_ANY` — the path name can be anything.
  - `_FTYPE_LINK` — reserved for the Process Manager.
  - `_FTYPE_MOUNT` — receive mount requests on the path (*path* must be NULL).
  - `_FTYPE_QUEUE` — reserved for a mqueue manager.
  - `_FTYPE_PIPE` — reserved for a pipe manager.

- `_FTYPE_SEM` — reserved for a semaphore manager.
- `_FTYPE_SHMEM` — reserved for a shared memory object.
- `_FTYPE_SOCKET` — reserved for a socket manager.
- `_FTYPE_SYMLINK` — reserved for the Process Manager.

*eflag*

Extended flags:

- `_IO_CONNECT_EFLAG_DIR` — the path referenced a directory.
- `_IO_CONNECT_EFLAG_DOT` — the last component of a path was `.` or `..` (i.e. the current or parent directory).

*chroot\_len*

The length of `chroot` in the returned path.

*umask*

One of:

- `S_IFBLK` — block special.
- `S_IFCHR` — character special.
- `S_IFDIR` — directory.
- `S_IFIFO` — FIFO special.
- `S_IFLNK` — symbolic link.
- `S_IFMT` — type of file.
- `S_IFNAM` — special named file.
- `S_IFREG` — regular.
- `S_IFSOCK` — socket.

*nentries*

If this member is zero, the path is a symbolic link.

*path\_len*

The length of the path including the terminating null character. If this member is zero, the path is null-terminated.

**Classification:**

QNX Neutrino

**See also:**

`_io_connect`, `_io_connect_fstype_reply`,  
`resmgr_connect_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.



## Synopsis:

```
#include <sys/ioctl.h>

int ioctl(int fd,
 int request,
 ...);
```

## Arguments:

- fd* An open file descriptor for the file or device that you want to manipulate.
- request* What you want to do to the file. The macros and definitions that you use in specifying a *request* are located in the file `<sys/ioctl.h>`.
- Additional arguments  
As required by the request.

## Library:

`libc`

## Description:

The *ioctl()* function manipulates the underlying parameters of files. In particular, it can be used to control many of the operating attributes of files (such as the attributes of terminals).

The *request* argument determines whether the subsequent arguments are an “in” or “out” parameter; it also specifies the size of the arguments in bytes.

## Returns:

A value based on the *request*, or -1 if an error occurs (*errno* is set).

**Errors:**

|        |                                                                                                                                                                                     |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF  | Invalid descriptor <i>fd</i> .                                                                                                                                                      |
| EINVAL | The <i>request</i> or optional variables aren't valid.                                                                                                                              |
| ENOTTY | The <i>fd</i> argument isn't associated with a character special device; the specified <i>request</i> doesn't apply to the kind of object that the descriptor <i>fd</i> references. |

**Classification:**

Standard Unix

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**Caveats:**

The *ioctl()* function is a Unix function that varies greatly from platform to platform.

**See also:**

*devctl()*

## Synopsis:

```
#include <sys/iomgr.h>

int iofdinfo(int filedes,
 unsigned flags,
 struct _fdinfo * info,
 char * path,
 int maxlen);
```

## Arguments:

- |                |                                                                                                                                                                                   |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>filedes</i> | A file descriptor for the connection that you want to query.                                                                                                                      |
| <i>flags</i>   | Specify <code>_FDINFO_FLAG_LOCALPATH</code> to return only the local path info (i.e. exclude the network path info).                                                              |
| <i>info</i>    | NULL, or a pointer to an <code>_fdinfo</code> structure that contains the connection information defined in <code>&lt;sys/iomgr.h&gt;</code> . Specify NULL if it's not required. |
| <i>path</i>    | A pointer to a buffer where the function can store the path associated with the file descriptor. Specify NULL if it's not required.                                               |
| <i>maxlen</i>  | The length of the buffer pointed to <i>path</i> .                                                                                                                                 |

## Library:

`libc`

## Description:

The *iofdinfo()* function retrieves the server's attribute information for the connection referred to by *filedes*.

**Returns:**

The length of the associated *filedes* pathname, or -1 if an error occurs (*errno* is set).

**Errors:**

- EFAULT      A fault occurred in a server's address space when it tried to access the caller's message buffers.
  
- EMSGSIZE    Insufficient space available in the server's buffer for the **fdinfo** data structure.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler      No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*iofunc\_fdinfo()*, *iofunc\_fdinfo\_default()*, *resmgr\_pathname()*

### **Synopsis:**

```
#include <sys/iofunc.h>

void iofunc_attr_init (iofunc_attr_t *attr,
 mode_t mode,
 iofunc_attr_t *dattr,
 struct _client_info *info);
```

### **Arguments:**

- attr*      A pointer to the `iofunc_attr_t` structure that you want to initialize.
- mode*      The type and access permissions that you want to use for the resource. For more information, see “Access permissions” in the documentation for `stat()`.
- dattr*      NULL, or a pointer to a `iofunc_attr_t` structure that you want to use to initialize the structure pointed to by *attr*.
- info*      NULL, or a pointer to a `_client_info` structure that contains the information about a client connection. For information about this structure, see `ConnectClientInfo()`.

### **Library:**

`libc`

### **Description:**

The `iofunc_attr_init()` function initializes the passed *attr* structure with the information derived from the optional *dattr*, the *mode*, and the user and group IDs from the optional *info* client information structure.

The *count*, *rcount*, *wcount*, *rlocks* and *wlocks* counters are reset to zero in the `iofunc_attr_t` structure that *attr* points to.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_lock()*, *iofunc\_attr\_t*, *iofunc\_attr\_unlock()*,  
*iofunc\_ocb\_attach()*, *iofunc\_ocb\_detach()*, *resmgr\_attach()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_attr_lock(iofunc_attr_t *attr);
```

### **Arguments:**

*attr* A pointer to the `iofunc_attr_t` structure that you want to lock.

### **Library:**

`libc`

### **Description:**

The `iofunc_attr_lock()` function locks the attribute structure that *attr* points to, preventing other threads in the resource manager from changing information.

Call this function (or `iofunc_attr_trylock()`) before you make any modifications to the attribute structure. After you're finished making modifications, call `iofunc_attr_unlock()` to release the lock.

Note that this is a *counting* locking mechanism. This means that a given thread can lock the attributes structure multiple times; it must then unlock the attributes structure a corresponding number of times in order to have the attributes structure considered unlocked. If another thread attempts to lock the structure while a thread has the structure locked, the other thread blocks.

### **Returns:**

EOK Success.

EAGAIN On the first use, all kernel mutex objects were in use.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_init()*, *iofunc\_attr\_t*, *iofunc\_attr\_trylock()*,  
*iofunc\_attr\_unlock()*

Writing a Resource Manager chapter of the *Programmer's Guide*.



**Synopsis:**

```

#include <sys/iofunc.h>

typedef struct _iofunc_attr {
 IOFUNC_MOUNT_T *mount; /* Used to find iofunc
 uint32_t flags; /* Dirty and invalid
 int32_t lock_tid; /* Thread that has at
 uint16_t lock_count; /* Lock count (0 == un
 uint16_t count; /* File use count */
 uint16_t rcount; /* File reader count
 uint16_t wcount; /* File writer count
 uint16_t rlocks; /* Number of read loc
 uint16_t wlocks; /* Number of write loc
 struct _iofunc_mmap_list *mmap_list; /* List of mmap ids *
 struct _iofunc_lock_list *lock_list; /* Lock lists */
 void *list; /* Reserved for futur
 uint32_t list_size; /* Size of reserved a

#if !defined(_IOFUNC_OFFSET_BITS) || _IOFUNC_OFFSET_BITS == 64
 #if _FILE_OFFSET_BITS - 0 == 64
 off_t nbytes; /* Always Number of by
 ino_t inode; /* mount point specif
 #else
 off64_t nbytes; /* Always Number of by
 ino64_t inode; /* mount point specif
 #endif
#elif _IOFUNC_OFFSET_BITS - 0 == 32
 #if !defined(_FILE_OFFSET_BITS) || _FILE_OFFSET_BITS == 32
 #if defined(__LITTLEENDIAN__)
 off_t nbytes; /* Always Number of by
 off_t nbytes_hi;
 ino_t inode; /* mount point specif
 ino_t inode_hi;
 #elif defined(__BIGENDIAN__)
 off_t nbytes_hi;
 off_t nbytes; /* Always Number of by
 ino_t inode_hi;
 ino_t inode; /* mount point specif
 #else
 #error endian not configured for system
 #endif
 #else
 #if defined(__LITTLEENDIAN__)

```

```

 int32_t nbytes; /* Always Number of bytes
 int32_t nbytes_hi;
 int32_t inode; /* mount point specific
 int32_t inode_hi;
 #elif defined(__BIGENDIAN__)
 int32_t nbytes_hi;
 int32_t nbytes; /* Always Number of bytes
 int32_t inode_hi;
 int32_t inode; /* mount point specific
 #else
 #error endian not configured for system
 #endif
#endif
#else
 #error _IOFUNC_OFFSET_BITS value is unsupported
#endif
 uid_t uid; /* User id */
 gid_t gid; /* Group id */
 time_t mtime; /* Modification time (write)
 time_t atime; /* Access time (read update)
 time_t ctime; /* Change time (write/change)
 mode_t mode; /* File mode (S_* from stat)
 nlink_t nlink; /* Number of links to the file
 dev_t rdev; /* dev num for CHR special
} iofunc_attr_t;

```

## Description:

The `iofunc_attr_t` structure describes the attributes of the device that's associated with a resource manager. The members include the following:

*mount*            A pointer a structure information about the mountpoint. By default, this structure is of type `iofunc_mount_t`, but you can specify your own structure by changing the `IOFUNC_MOUNT_T` manifest.

*flags*

Flags that your resource manager can set to indicate the state of the device. This member is a combination of the following flags:

## IOFUNC\_ATTR\_ETIME

The access time is no longer valid. Typically set on a read from the resource.

## IOFUNC\_ATTR\_CTIME

The change of status time is no longer valid. Typically set on a file info change.

## IOFUNC\_ATTR\_DIRTY\_NLINK

The number of links has changed.

## IOFUNC\_ATTR\_DIRTY\_MODE

The mode has changed.

## IOFUNC\_ATTR\_DIRTY\_OWNER

The uid or the gid has changed.

## IOFUNC\_ATTR\_DIRTY\_RDEV

The *rdev* member has changed, e.g. *mknod()*.

## IOFUNC\_ATTR\_DIRTY\_SIZE

The size has changed.

## IOFUNC\_ATTR\_DIRTY\_TIME

One or more of *mtime*, *atime*, or *ctime* has changed.

## IOFUNC\_ATTR\_MTIME

The modification time is no longer valid. Typically set on a write to the resource.

In addition to the above, your resource manager can use in any way the bits in the range defined by IOFUNC\_ATTR\_PRIVATE (see `<sys/iofunc.h>`).

*lock\_tid*

The ID of the thread that has locked the attributes. To support multiple threads in your resource manager,

you'll need to lock the attribute structure so that only one thread at a time is allowed to change it.

The resource manager layer automatically locks the attribute (using for you when certain handler functions are called (i.e. IO\_\*).

*lock\_count* The number of times the thread has locked the attribute structure. You can lock the attributes by calling *iofunc\_attr\_lock()* or *iofunc\_attr\_trylock()*; unlock them by calling *iofunc\_attr\_unlock()*




---

A thread must unlock the attributes as many times as it locked them.

---

*count* The number of OCBs using this attribute in any manner. When this count is zero, no one is using this attribute.

*rcount* The number of OCBs using this attribute for reading.

*wcount* The number of OCBs using this attribute for writing.

*rlocks* The number of read locks currently registered on the attribute.

*wlocks* The number of write locks currently registered on the attribute.

*mmap\_list* and *lock\_list*

To manage their particular functionality on the resource, the *mmap\_list* member is used by *iofunc\_mmap()* and *iofunc\_mmap\_default()*; the *lock\_list* member is used by *iofunc\_lock\_default()*. Generally, you shouldn't need to modify or examine these members.

*list* Reserved for future use.

*list\_size* Size of reserved area; reserved for future use.

- nbytes* The number of bytes in the resource; your resource manager can change this value.
- For a file, this would contain the file's size. For special devices (e.g. `/dev/null`) that don't support `lseek()` or have a radically different interpretation for `lseek()`, this field isn't used (because you wouldn't use any of the helper functions, but would supply your own instead.) In these cases, we recommend that you set this field to zero, unless there's a meaningful interpretation that you care to put to it.
- inode* This is a mountpoint-specific inode that must be unique per mountpoint. You can specify your own value, or 0 to have the Process manager fill it in for you. For filesystem type of applications, this may correspond to some on-disk structure. In any case, the interpretation of this field is up to you.
- uid* and *gid* The user ID and group ID of the owner of this resource. These fields are updated automatically by the `chown()` helper functions (e.g. `iofunc_chown_default()`) and are referenced in conjunction with the *mode* member for access-granting purposes by the `open()` help functions (e.g. `iofunc_open_default()`).
- mtime*, *atime*, and *ctime*
- POSIX time members:
- *mtime* — modification time (`write()` updates this).
  - *atime* — access time (`read()` updates this).
  - *ctime* — change of status time (`write()`, `chmod()` and `chown()` update this).



One or more of the three time members may be *invalidated* as a result of calling an iofunc-layer function. To see if a time member is invalid, check the *flags* member. This is to avoid having each and every I/O message handler go to the kernel and request the current time of day, just to fill in the attribute structure's time member(s).

---

To fill the members with the correct time, call *iofunc\_time\_update()*.

|              |                                                                                                                                                                                                                   |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>mode</i>  | The resource's mode (e.g. type, permissions). Valid modes may be selected from the S_* series of constants in <code>&lt;sys/stat.h&gt;</code> ; see "Access permissions" in the documentation for <i>stat()</i> . |
| <i>nlink</i> | The number of links to this particular name; your resource manager can modify this member. For names that represent a directory, this value must be greater than 2.                                               |
| <i>rdev</i>  | The device number for a character special device and the <code>rdev</code> number for a named special device.                                                                                                     |

## Classification:

QNX Neutrino

## See also:

*iofunc\_attr\_lock()*, *iofunc\_attr\_trylock()*, *iofunc\_attr\_unlock()*,  
*iofunc\_lock\_default()*, *iofunc\_mmap()*, *iofunc\_mmap\_default()*,  
**iofunc\_ocb\_t**, *iofunc\_time\_update()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_attr_trylock(iofunc_attr_t *attr);
```

**Arguments:**

*attr* A pointer to the `iofunc_attr_t` structure that you want to lock.

**Library:**

`libc`

**Description:**

The `iofunc_attr_trylock()` function attempts to lock the attribute structure *attr*, preventing other threads in the resource manager from changing information. If it can't lock *attr* immediately, it returns `EBUSY`.

Call this function (or `iofunc_attr_lock()`) before you make any modifications to the attribute structure. After you're finished making modifications, call `iofunc_attr_unlock()` to release the lock.

Note that this is a *counting* locking mechanism. This means that a given thread can lock the attributes structure multiple times; it must then unlock the attributes structure a corresponding number of times in order to have the attributes structure considered unlocked. If another thread attempts to lock the structure while a thread has the structure locked, the other thread will block.

**Returns:**

|                     |                                                              |
|---------------------|--------------------------------------------------------------|
| <code>EOK</code>    | Success.                                                     |
| <code>EBUSY</code>  | The calling thread couldn't lock the attributes immediately. |
| <code>EAGAIN</code> | On the first use, all kernel mutex objects were in use.      |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_init()*, *iofunc\_attr\_lock()*, *iofunc\_attr\_t*,  
*iofunc\_attr\_unlock()*

Writing a Resource Manager chapter of the *Programmer's Guide*.



### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_attr_unlock(iofunc_attr_t *attr);
```

### **Arguments:**

*attr* A pointer to the `iofunc_attr_t` structure that you want to unlock.

### **Library:**

`libc`

### **Description:**

The `iofunc_attr_unlock()` function unlocks the attribute structure *attr*, allowing other threads in the resource manager to change information.

Use this function in conjunction with `iofunc_attr_lock()` or `iofunc_attr_trylock()`; call `iofunc_attr_unlock()` after you've made modifications to the attribute structure. You must unlock the structure as many times as you locked it.

### **Returns:**

EOK Success.

EAGAIN On the first use, all kernel mutex objects were in use.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point No  
*continued...*

## **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

## **See also:**

*iofunc\_attr\_init()*, *iofunc\_attr\_lock()*, *iofunc\_attr\_t*,  
*iofunc\_attr\_trylock()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_check_access(
 resmgr_context_t *ctp,
 const iofunc_attr_t *attr,
 mode_t checkmode,
 const struct _client_info *info);
```

**Arguments:**

|                  |                                                                                                                                                                                                                                                                    |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>       | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions.                                                                                                                       |
| <i>attr</i>      | A pointer to the <code>iofunc_attr_t</code> structure that defines the characteristics of the device that's associated with the resource manager.                                                                                                                  |
| <i>checkmode</i> | The type and access permissions that you want to check for the resource. For more information, see below.                                                                                                                                                          |
| <i>info</i>      | A pointer to a <code>_client_info</code> structure that contains the information about a client connection. For information about this structure, see <code>ConnectClientInfo()</code> . You can get this structure by calling <code>iofunc_client_info()</code> . |

**Library:**

`libc`

**Description:**

The `iofunc_check_access()` function verifies that the client is allowed access to the resource, as specified by a combination of who the client is (*info*), and the resource attributes `attr->mode`, `attr->uid` and `attr->gid`. Access is tested based upon the *checkmode* parameter.

The *checkmode* parameter determines which checks are done. It's a bitwise OR of the following constants:

- S\_ISUID Verifies that the effective user ID of the client is equal to the user ID specified by the *attr->uid* member.
- S\_ISGID Verifies that the effective group ID or one of the supplementary group IDs of the client is equal to the group ID specified by the *attr->gid* member.
- S\_IREAD Verifies that the client has READ access to the resource as specified by *attr->mode*.  
  
If the client's effective user ID matches that of *attr->uid*, then the permission check is made against the owner permission field of *attr->mode* (mask 0700 octal).  
  
If the client's effective user ID doesn't match that of *attr->uid*, then if the client's effective group ID matches that of *attr->gid*, or one of the client's supplementary group IDs matches *attr->gid*, the check is made against the group permission field of *attr->mode* (mask 0070 octal).  
  
If none of the group fields match, the check is made against the other permission field of *attr->mode* (mask 0007 octal).
- S\_IWRITE Same as S\_IREAD, except WRITE access is tested.
- S\_IEXEC Same as S\_IREAD, except EXECUTE access is tested. Note that since most resource managers don't actually execute code, the execute access is typically used in its directory sense, i.e. to test for directory accessibility, rather than execute access.

The S\_ISUID and S\_ISGID flags are mutually exclusive, that is, you may specify at most one of them. In conjunction with the S\_ISUID and S\_ISGID flags, you may specify zero or more of the S\_IREAD,

S\_IWRITE, and S\_IEXEC flags. If no flags are specified, the permission checks are performed for privileged (**root**) access.

Here's some pseudo-code to try to explain this:

```
if superuser:
 return EOK

if S_ISUID and effective user ID == file user ID:
 return EOK

if S_ISGID and effective group ID == file group ID:
 return EOK

if S_IREAD or S_IWRITE or S_IEXEC:
 if caller's user ID == effective user ID:
 if all permissions are set in file's owner mode bits:
 return EOK
 else:
 return EACCESS

 if (caller's group ID or supplementary group IDs) ==
 effective group ID:
 if all permissions are set in file's group mode bits:
 return EOK
 else:
 return EACCESS

 if all permissions are set in file's other mode bits:
 return EOK
 else:
 return EACCESS

return EPERM
```

## Returns:

|        |                                                          |
|--------|----------------------------------------------------------|
| EACCES | The client doesn't have permissions to do the operation. |
| ENOSYS | NULL was passed for <i>info</i> structure.               |
| EOK    | Successful completion.                                   |
| EPERM  | The group ID or owner ID didn't match.                   |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_client\_info()*, *iofunc\_open()*, *iofunc\_read\_verify()*,  
*iofunc\_write\_verify()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_chmod (resmgr_context_t *ctp,
 io_chmod_t *msg,
 iofunc_ocb_t *ocb,
 iofunc_attr_t *attr);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_chmod_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

## Library:

`libc`

## Description:

The `iofunc_chmod()` helper function implements POSIX semantics for the client's `chmod()` call, which is received as an `_IO_CHMOD` message by the resource manager.

The `iofunc_chmod()` function verifies that the client has the necessary permissions to effect a `chmod()` on the attribute. If so, the `chmod()` is performed, modifying elements of the `ocb->attr` structure. This

function takes care of updating the IOFUNC\_ATTR\_CTIME, IOFUNC\_ATTR\_DIRTY\_TIME, and IOFUNC\_ATTR\_DIRTY\_MODE bits in *ocb->attr->flags*. You can use *iofunc\_time\_update()*, to update the appropriate time fields in *ocb->attr*.

You can use *iofunc\_chmod()*, for example, in a filesystem manager, where an `_IO_CHMOD` message was received, and the filesystem code must now write the values to the medium. The filesystem code may wish to block the client thread until the data was actually written to the medium. Contrast this scenario to the behavior of *iofunc\_chmod\_default()*, which calls this routine, and replies to the client thread.

## **io\_chmod\_t structure**

The `io_chmod_t` structure holds the `_IO_CHMOD` message received by the resource manager:

```
struct _io_chmod {
 uint16_t type;
 uint16_t combine_len;
 mode_t mode;
};

typedef union {
 struct _io_chmod i;
} io_chmod_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_chmod` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_CHMOD</code> .                                                                                                                                                                                            |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see "Combine messages" in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |



*mode*            The new mode. For more information, see “Access permissions” in the documentation for *stat()*.

## Returns:

EOK            Successful completion.

EROFS         An attempt was made to chmod on a read-only filesystem.

EACCES        The client doesn't have permissions to do the operation.

EPERM         The group ID or owner ID didn't match.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_attr\_t*, *iofunc\_chmod\_default()*, *iofunc\_ocb\_t*,  
*iofunc\_time\_update()*, *resmgr\_context\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_chmod\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_CHMOD messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_chmod_default(resmgr_context_t *ctp,
 io_chmod_t *msg,
 iofunc_ocr_t *ocr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_chmod_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_chmod()`.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_chmod_default()` function implements POSIX semantics for the client's `chmod()` call, which is received as an `_IO_CHMOD` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `chmod` position, or you can call `iofunc_func_init()` to initialize all the functions to their default values.

The `iofunc_chmod_default()` function calls `iofunc_chmod()` to do the actual work, and (if installed in the `io_funcs` table) issues the reply back to the client.

## Returns:

|        |                                                                  |
|--------|------------------------------------------------------------------|
| EOK    | Successful completion.                                           |
| EROFS  | An attempt was made to <i>chmod()</i> on a read-only filesystem. |
| EACCES | The client doesn't have permissions to do the operation.         |
| EPERM  | The group ID or owner ID didn't match.                           |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_chmod()*, *iofunc\_func\_init()*, *iofunc\_ocb\_t*,  
*iofunc\_time\_update()*, *resmgr\_attach()*, *resmgr\_context\_t*,  
*resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_chown()***

© 2004, QNX Software Systems Ltd.

*Handle an \_IO\_CHOWN message*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_chown (resmgr_context_t *ctp,
 io_chown_t *msg,
 iofunc_ocr_t *ocr,
 iofunc_attr_t *attr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_chown_t` structure that contains the message that the resource manager received; see below.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The `iofunc_chown()` helper function implements POSIX semantics for the client's `chown()` call, which is received as an `_IO_CHOWN` message by the resource manager.

The `iofunc_chown()` function verifies that the client has the necessary permissions to effect a `chown` on the attribute. If so, the `chown` is performed, modifying elements of the `ocr->attr` structure. As per

POSIX 1003.1, if the client isn't `root`, *iofunc\_chown()* clears the set-user-id and set-group-id bits in the *ocb->attr->mode* member.

This function takes care of updating the `IOFUNC_ATTR_CTIME`, `IOFUNC_ATTR_DIRTY_TIME`, and `IOFUNC_ATTR_DIRTY_MODE` bits in *ocb->attr->flags*. You can use *iofunc\_time\_update()*, to update the appropriate time fields in *ocb->attr*.

### **io\_chown\_t structure**

The `io_chown_t` structure holds the `_IO_CHOWN` message received by the resource manager:

```
struct _io_chown {
 uint16_t type;
 uint16_t combine_len;
 int32_t gid;
 int32_t uid;
};

typedef union {
 struct _io_chown i;
} io_chown_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_chown` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_CHOWN</code> .                                                                                                                                                                                            |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see "Combine messages" in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |
| <i>gid</i>         | The new group ID.                                                                                                                                                                                                   |
| <i>uid</i>         | The new user ID.                                                                                                                                                                                                    |

**Returns:**

|        |                                                          |
|--------|----------------------------------------------------------|
| EOK    | Successful completion.                                   |
| EROFS  | An attempt was made to chown on a read-only filesystem.  |
| EACCES | The client doesn't have permissions to do the operation. |
| EPERM  | The group ID or owner ID didn't match.                   |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_attr_t`, `iofunc_chmod()`, `iofunc_chown_default()`,  
`iofunc_ocb_t`, `iofunc_time_update()`, `resmgr_attach()`,  
`resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_chown_default(resmgr_context_t *ctp,
 io_chown_t *msg,
 iofunc_ocb_t *ocb);
```

**Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_chown_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_chown()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

**Library:**

`libc`

**Description:**

The `iofunc_chown_default()` function implements POSIX semantics for the client's `chown()` call, which is received as an `_IO_CHOWN` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `chown` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_chown_default()` function calls `iofunc_chown()` to do the actual work, and (if installed in the `io_funcs` table) issues the reply back to the client.

**Returns:**

|        |                                                          |
|--------|----------------------------------------------------------|
| EOK    | Successful completion.                                   |
| EROFS  | An attempt was made to chown on a read-only filesystem.  |
| EACCES | The client doesn't have permissions to do the operation. |
| EPERM  | The group ID or owner ID didn't match.                   |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_chown()*, *iofunc\_func\_init()*, `iofunc_ocb_t`,  
*iofunc\_time\_update()*, *resmgr\_attach()*, `resmgr_context_t`,  
`resmgr_io_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.



### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_client_info (resmgr_context_t * ctp,
 int ioflag,
 struct _client_info * info);
```

### **Arguments:**

- ctp* A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions.
- ioflag* Zero, or the constant **O\_REALIDS**. This argument is passed in the **\_JO\_OPEN** message during an open request. If **O\_REALIDS** is specified, *iofunc\_client\_info()* swaps the real and effective values of the user and group IDs before returning. This is a QNX Neutrino extension, to swap real and effective user and group IDs in an atomic operation.
- info* A pointer to a **\_client\_info** structure that the function fills with information about a client connection. For information about this structure, see *ConnectClientInfo()*.

### **Library:**

**libc**

### **Description:**

The *iofunc\_client\_info()* function fetches the *info* structure for the client. It calls *ConnectClientInfo()* to gather the information, based on the server connection ID found in *ctp->info.scoid*.

**Returns:**

- EFAULT     A fault occurred when the kernel tried to access the *info* buffer provided.
- EINVAL     The client process is no longer valid.
- EOK        Successful completion.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*ConnectClientInfo()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_close_dup(resmgr_context_t* ctp,
 io_close_t* msg,
 iofunc_ocr_t* ocb,
 iofunc_attr_t* attr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_close_t` structure that contains the message that the resource manager received; see below.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The `iofunc_close_dup()` helper function handles a `_JO_CLOSE` message. This function frees all locks allocated for the client process on the file descriptor and performs any POSIX-related cleanup required when a duplicated `ocr` is detached.

**io\_close\_t structure**

The `io_close_t` structure holds the `_IO_CLOSE` message received by the resource manager:

```
struct _io_close {
 uint16_t type;
 uint16_t combine_len;
};

typedef union {
 struct _io_close i;
} io_close_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_close` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_CLOSE</code> .                                                                                                                                                                                            |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see "Combine messages" in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |

**Returns:**

|               |                    |
|---------------|--------------------|
| EOK           | Success.           |
| Anything else | An error occurred. |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_close\_dup\_default()*, *iofunc\_close\_ocb()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_close\_dup\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_CLOSE messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_close_dup_default(
 resmgr_context_t *ctp,
 io_close_t *msg,
 iofunc_ocb_t *ocb);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_close_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_close_dup()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_close_dup_default()` function implements default actions for the `_IO_CLOSE` message. This function simply calls `iofunc_close_dup()`, which does the actual work.

You can place `iofunc_close_dup_default()` directly into the `io_funcs` table passed to `resmgr_attach()`, at the `close_dup` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.



---

If your resource manager uses *iofunc\_lock\_default()*, you *must* use both this function (*iofunc\_close\_dup\_default()*) and *iofunc\_unblock\_default()*, as they provide necessary ancillary functionality for managing file locks. This is because file locks are owned by the process, and aren't inherited by the child process.

---

**Returns:**

EOK      Successful completion.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_close\_dup()*, *iofunc\_func\_init()*, *iofunc\_ocb\_t*,  
*iofunc\_time\_update()*, *resmgr\_attach()*, *resmgr\_context\_t*,  
*resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_close\_ocb()***

© 2004, QNX Software Systems Ltd.

*Return the memory allocated for an OCB*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_close_ocb(resmgr_context_t* ctp,
 iofunc_ocb_t* ocb,
 iofunc_attr_t* attr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The `iofunc_close_ocb()` function detaches the OCB specified by *ocb*, and releases the memory associated with it.





---

This function assumes that *ocb* points to an `iofunc_ocb_t`. If you encapsulate `iofunc_ocb_t` in your own OCB it must be the first field of your OCB; otherwise, you can't call this function. If you provide an *ocb\_free()* function in the mount structure then it's called at this point. This means that at least the `iofunc_ocb_t` portion of your OCB is no longer valid after *iofunc\_close\_ocb()* returns.

---

The *iofunc\_close\_ocb()* function calls *iofunc\_ocb\_detach()* on your behalf.

### Returns:

EOK      Success.

### Classification:

QNX Neutrino

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

`iofunc_attr_t`, *iofunc\_close\_dup()*, *iofunc\_close\_ocb\_default()*, *iofunc\_ocb\_free()*, `iofunc_ocb_t`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_close\_ocr\_default()***

© 2004, QNX Software Systems Ltd.

*Return the memory allocated for an OCB*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_close_ocr_default(resmgr_context_t* ctp,
 void* reserved,
 iofunc_ocr_t* ocb);
```

### **Arguments:**

|                 |                                                                                                                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <b>resmgr_context_t</b> structure that the resource-manager library uses to pass context information between functions. |
| <i>reserved</i> | This argument must be passed as NULL.                                                                                                  |
| <i>ocr</i>      | A pointer to the <b>iofunc_ocr_t</b> structure for the Open Control Block that was created when the client opened the resource.        |

### **Library:**

**libc**

### **Description:**

The *iofunc\_close\_ocr\_default()* function detaches the OCB specified by *ocr*, and releases the memory associated with it.

You can place this function directly into the *io\_funcs* table passed to *resmgr\_attach()*, at the *close\_ocr* position, or you can call *iofunc\_func\_init()* to initialize all of the functions to their default values.



---

This function assumes that *ocr* points to an `iofunc_ocr_t`. If you encapsulate `iofunc_ocr_t` in your own OCB, it must be the first field of your OCB; otherwise, you can't call this function. If you provide an *ocr\_free()* function in the mount structure, it's called at this point. This means that at least the `iofunc_ocr_t` portion of your OCB is no longer valid after *iofunc\_close\_ocr()* returns.

---

The *iofunc\_close\_ocr\_default()* function calls *iofunc\_close\_ocr()*.

## Returns:

EOK      Success.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_close\_ocr()*, *iofunc\_func\_init()*, `iofunc_ocr_t`,  
*iofunc\_time\_update()*, *resmgr\_attach()*, `resmgr_context_t`,  
`resmgr_io_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_devctl()***

© 2004, QNX Software Systems Ltd.

*Handle an \_IO\_DEVCTL message*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_devctl(resmgr_context_t *ctp,
 io_devctl_t *msg,
 iofunc_ocb_t *ocb,
 iofunc_attr_t *attr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_devctl_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The `iofunc_devctl()` helper function implements POSIX semantics for the client's `devctl()` call, which is received as an `_IO_DEVCTL` message by the resource manager. This function handles the `DCMD_ALL*` functionality.

This function handles at least the following device control messages:

`DCMD_ALL_GETFLAGS`

Implements the functionality of the `fcntl()` get-flags command.

**DCMD\_ALL\_SETFLAGS**

Implements the functionality of the *fcntl()* set-flags command.

**DCMD\_ALL\_GETMOUNTFLAGS**

Returns the mount flag (*mount->flags*) for a resource that has a mount structure defined, else returns a mount flag of zero.

The supported mount flags (bitmask values) for DCMD\_ALL\_GETMOUNTFLAGS include:

**\_MOUNT\_READONLY**

Read only.

**\_MOUNT\_NOEXEC**

Can't exec from filesystem.

**\_MOUNT\_NOSUID**

Don't honor setuid bits on filesystem.

Any other device control messages return ENOTTY.

**io\_devctl\_t structure**

The *io\_devctl\_t* structure holds the *\_IO\_* message received by the resource manager:

```

struct _io_devctl {
 uint16_t type;
 uint16_t combine_len;
 int32_t dcmd;
 int32_t nbytes;
 int32_t zero;
 /* char data[nbytes]; */
};

struct _io_devctl_reply {
 uint32_t zero;
 int32_t ret_val;
 int32_t nbytes;
 int32_t zero2;
 /* char data[nbytes]; */
};

```

```
};

typedef union {
 struct _io_devctl i;
 struct _io_devctl_reply o;
} io_devctl_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The *i* member is a structure of type `_io_devctl` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_DEVCTL</code> .                                                                                                                                                                                           |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> . |
| <i>dcmd</i>        | The device-control command to execute.                                                                                                                                                                              |
| <i>nbytes</i>      | The number of bytes of data being passed with the command.                                                                                                                                                          |

The commented-out declaration for *data* indicates that *nbytes* bytes of data immediately follow the `_io_devctl` structure.

The `_DEVCTL_DATA()` macro gets a pointer to the data that follows the message. Call it like this:

```
data = _DEVCTL_DATA (msg->i);
```

The *o* member of the `io_devctl_t` message is a structure of type `_io_devctl_reply` that contains the following members:

|                |                                             |
|----------------|---------------------------------------------|
| <i>ret_val</i> | The value returned by the command.          |
| <i>nbytes</i>  | The number of bytes of data being returned. |

The commented-out declaration for *data* indicates that *nbytes* bytes of data immediately follow the `_io_devctl_reply` structure.

## Returns:

|        |                                                                                                                                   |
|--------|-----------------------------------------------------------------------------------------------------------------------------------|
| EOK    | Successful completion.                                                                                                            |
| EINVAL | An attempt to set the flags for a resource that is synchronized, with no mount structure defined, or no synchronized I/O defined. |
| ENOTTY | An unsupported device control message was decoded.                                                                                |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

`fcntl()`, `iofunc_attr_t`, `iofunc_devctl_default()`, `iofunc_ocb_t`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_devctl\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_DEVCTL messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_devctl_default(resmgr_context_t *ctp,
 io_devctl_t *msg,
 iofunc_ocb_t *ocb);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_devctl_t` structure that contains the message that the resource manager received. For more information, see the documentation for *iofunc\_devctl()*.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The *iofunc\_devctl\_default()* function implements POSIX semantics for the client's *devctl()* call, which is received as an `_IO_DEVCTL` message by the resource manager.

You can place this function directly into the *io\_funcs* table passed to *resmgr\_attach()*, at the *devctl* position, or you can call *iofunc\_func\_init()* to initialize all of the functions to their default values.

The *iofunc\_devctl\_default()* function calls *iofunc\_devctl()*, to do the actual work, and (if installed in the *io\_funcs* table) issues the reply back to the client.



**Returns:**

|                  |                                                                                                                                   |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| EOK              | Successful completion.                                                                                                            |
| EINVAL           | An attempt to set the flags for a resource that is synchronized, with no mount structure defined, or no synchronized I/O defined. |
| ENOTTY           | An unsupported device control message was decoded.                                                                                |
| ._RESMGR_DEFAULT | An supported device control message that isn't a known DCMD_ALL_* command was decoded.                                            |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_devctl()*, *iofunc\_func\_init()*, **iofunc\_ocb\_t**,  
*iofunc\_time\_update()*, *resmgr\_attach()*, **resmgr\_context\_t**,  
**resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_fdinfo()***

© 2004, QNX Software Systems Ltd.

*Handle an IO\_FDINFO message*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_fdinfo(resmgr_context_t * ctp,
 iofunc_ocb_t * ocb,
 iofunc_attr_t * attr,
 struct _fdinfo * info);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* NULL, or a pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.
- info* A pointer to a `_fdinfo` structure that the function fills with the information. This structure is defined in `<sys/iomgr.h>` as:

```
struct _fdinfo {
 uint32_t mode; /* File mode */
 uint32_t ioflag; /* Current io flags */
 uint64_t offset; /* Current seek position */
 uint64_t size; /* Current size of file */
 uint32_t flags; /* _FDINFO_* */
 uint16_t sflag; /* Share flags */
 uint16_t count; /* File use count */
 uint16_t rcount; /* File reader count */
 uint16_t wcount; /* File writer count */
 uint16_t rlocks; /* Number of read locks */
 uint16_t wlocks; /* Number of write locks */
 uint32_t zero[6];
};
```

The `_fdinfo` structure is included in the reply part of a `io_fdinfo_t` structure; for more information, see the documentation for `iofunc_fdinfo_default()`.

## Library:

`libc`

## Description:

The `iofunc_fdinfo()` helper function provides the implementation for the client's `iofdinfo()` call, which is received as an `_JO_FDINFO` message by the resource manager.

The `iofunc_fdinfo()` function transfers the appropriate fields from the `ocb` and `attr` structures to the `info` structure. If `attr` is `NULL`, then the `attr` information comes from the structure pointed to by `ocb->attr`.

## Returns:

`EOK` Successful completion.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofdinfo()*, `iofunc_attr_t`, *iofunc\_fdinfo\_default()*,  
`iofunc_ocb_t`, `resmgr_context_t`, *resmgr\_pathname()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iomgr.h>

int iofunc_fdinfo_default(resmgr_context_t * ctp,
 io_fdinfo_t * msg,
 iofunc_ocb_t * ocb);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_fdinfo_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

## Library:

`libc`

## Description:

The `iofunc_fdinfo_default()` function provides the default handler for the client's `iofdinfo()` call, which is received as an `_IO_FDINFO` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `fdinfo` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_fdinfo_default()` function calls `iofunc_fdinfo()` and `resmgr_pathname()` to do the actual write, and (if installed in the `io_funcs` table) replies back to the client.

**io\_fdinfo\_t structure**

The `io_fdinfo_t` structure holds the `_IO_FDINFO` message received by the resource manager:

```
struct _io_fdinfo {
 uint16_t type;
 uint16_t combine_len;
 uint32_t flags;
 int32_t path_len;
 uint32_t reserved;
};

struct _io_fdinfo_reply {
 uint32_t zero[2];
 struct _fdinfo info;
 /* char path[path_len + 1]; */
};

typedef union {
 struct _io_fdinfo i;
 struct _io_fdinfo_reply o;
} io_fdinfo_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The *i* member is a structure of type `_io_fdinfo` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_FDINFO</code> .                                                                                                                                                                                           |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> . |
| <i>flags</i>       | Specify <code>_FDINFO_FLAG_LOCALPATH</code> to return only the local path info (i.e. exclude the network path info).                                                                                                |
| <i>path_len</i>    | The size of the path reply buffers that follow the reply.                                                                                                                                                           |

The *o* member is a structure of type `_io_fdinfo_reply` that contains the following members:

*info* A `_fdinfo` structure that's defined (in `<sys/iomgr.h>`) as:

```

struct _fdinfo {
 uint32_t mode; /* File mode */
 uint32_t ioflag; /* Current io flags */
 uint64_t offset; /* Current seek position */
 uint64_t size; /* Current size of file */
 uint32_t flags; /* _FDINFO_* */
 uint16_t sflag; /* Share flags */
 uint16_t count; /* File use count */
 uint16_t rcount; /* File reader count */
 uint16_t wcount; /* File writer count */
 uint16_t rlocks; /* Number of read locks */
 uint16_t wlocks; /* Number of write locks */
 uint32_t zero[6];
};

```

The commented-out declaration for *path* indicates that *path* *len* + 1 bytes of data immediately follow the `_io_fdinfo_reply` structure.

## Returns:

The length of the path, or -1 if an error occurs (*errno* is set).

## Errors:

EMSGSIZE Insufficient space available in the server's buffer to receive the entire message.

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofdinfo()*, *iofunc\_fdinfo\_default()*, *iofunc\_func\_init()*, *iofunc\_ocb\_t*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*,  
*RESMGR\_NPARTS()*, *resmgr\_pathname()*

Writing a Resource Manager chapter of the *Programmer's Guide*.



### Synopsis:

```
#include <sys/iofunc.h>

void iofunc_func_init(
 unsigned nconnect,
 resmgr_connect_funcs_t *connect,
 unsigned nio,
 resmgr_io_funcs_t *io);
```

### Arguments:

- |                 |                                                                                                                                                     |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>nconnect</i> | The number of entries in the <i>connect</i> table that you want to fill. Typically, you pass <code>_RESMGR_CONNECT_NFUNCS</code> for this argument. |
| <i>connect</i>  | A pointer to a <code>resmgr_connect_funcs_t</code> structure that you want to fill with the default connect functions.                              |
| <i>nio</i>      | The number of entries in the <i>io</i> table that you want to fill. Typically, you pass <code>_RESMGR_IO_NFUNCS</code> for this argument.           |
| <i>io</i>       | A pointer to a <code>resmgr_io_funcs_t</code> structure that you want to fill with the default I/O functions.                                       |

### Library:

`libc`

### Description:

The *iofunc\_func\_init()* function initializes the passed *connect* and *io* structures with the POSIX-layer default functions. For information about the default functions, see `resmgr_connect_funcs_t` and `resmgr_io_funcs_t`.

The *nconnect* and *nio* arguments indicate how many entries this function should fill. This is in place to support forward compatibility.

## Examples:

Fill a connect and I/O function table with the POSIX-layer defaults:

```
#include <sys/iofunc.h>

static resmgr_connect_funcs_t my_connect_functions;
static resmgr_io_funcs_t my_io_functions;

int main (int argc, char **argv)
{
 ...
 iofunc_func_init (_RESMGR_CONNECT_NFUNCS, &my_connect_functions,
 _RESMGR_IO_NFUNCS, &my_io_functions);

 /*
 * At this point, the defaults have been filled in.
 * You may now override some of the default functions with
 * functions that you have written:
 */

 my_io_functions.io_read = my_io_read;
 ...
}
```

The above example initializes your connect and I/O function structures (*my\_connect\_functions* and *my\_io\_functions*) with the POSIX-layer defaults. If you didn't override any of the functions, your resource manager would behave like `/dev/null` — any data written to it would be discarded, and an attempt to read data from it would immediately return an EOF.

Since this isn't desirable in most cases, you'll often provide functionality for some functions, such as reading, writing, and device control to your device. In the example above, we've explicitly supplied our own handler for reading from the device, via a function called *my\_io\_read()*.

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_init()*, *resmgr\_attach()*, `resmgr_connect_funcs_t`,  
`resmgr_io_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_link()***

© 2004, QNX Software Systems Ltd.

*Link two directories*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_link(resmgr_context_t* ctp,
 io_link_t* msg,
 iofunc_attr_t* attr,
 iofunc_attr_t* dattr,
 struct _client_info* info);
```

### **Arguments:**

- ctp* A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the **io\_link\_t** structure that contains the message that the resource manager received; see below.
- attr* A pointer to the **iofunc\_attr\_t** structure that describes the characteristics of the resource.
- dattr* NULL, or a pointer to the **iofunc\_attr\_t** structure that describes the characteristics of the parent directory.
- info* NULL, or a pointer to a **\_client\_info** structure that contains information about the client. For information about this structure, see *ConnectClientInfo()*.

### **Library:**

**libc**

### **Description:**

The *iofunc\_link()* helper function links directory *attr* to *dattr* for context *ctp*. It's similar to the *iofunc\_open()* function:

The *iofunc\_link()* function checks to see if the client (described by the optional *info* structure) has access to open the resource (name passed

in the *msg* structure). The *attr* structure describes the resource's attributes, and the optional *dattr* structure defines the attributes of the parent directory (i.e. if *dattr* isn't NULL, it implies that the resource identified by *attr* is being created within the directory specified by *dattr*).

You can pass the *info* argument as NULL, in which case *iofunc\_link()* obtains the client information itself via a call to *iofunc\_client\_info()*. It is, of course, more efficient to get the client info once, rather than calling this function with NULL every time.

If you pass NULL in *info*, the function returns information about a client's connection in *info*, and an error constant.

### **io\_link\_t structure**

The *io\_link\_t* structure holds the `_IO.CONNECT` message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_ftype_reply ftype_reply;
} io_link_t;
```

This message structure is a union of an input message (coming to the resource manager), `_io_connect`, and two possible output or reply messages (going back to the client):

- `_io_connect_link_reply` if the reply is redirecting the client to another resource

Or:

- `_io_connect_ftype_reply` if the reply consists of a status and a file type.

The reply includes the following additional information:

```
struct _io_resmgr_link_extra {
 uint32_t nd;
 int32_t pid;
 int32_t chid;
 uint32_t handle;
```

```
uint32_t flags;
uint32_t file_type;
uint32_t reserved[2];
};

typedef union _io_link_extra {
 struct _msg_info info; /* EXTRA_LINK (from client) */
 void *ocb; /* EXTRA_LINK (from resmgr functions) */
 char path[1]; /* EXTRA_SYMLINK */
 struct _io_resmgr_link_extra resmgr; /* EXTRA_RESMGR_LINK */
} io_link_extra_t;
```

*info* A pointer to a `_msg_info` structure.

## Returns:

|           |                                                                          |
|-----------|--------------------------------------------------------------------------|
| EOK       | Success.                                                                 |
| EBADFSYS. | NULL was passed in <i>attr</i> and <i>dattr</i> .                        |
| EFAULT    | A fault occurred when the kernel tried to access the <i>info</i> buffer. |
| EINVAL    | The client process is no longer valid.                                   |
| ENOSYS    | NULL was passed in <i>info</i> .                                         |
| EPERM     | The group ID or owner ID didn't match.                                   |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectClientInfo()*, `iofunc_attr_t`, *iofunc\_client\_info()*,  
*iofunc\_open()*, `_msg_info`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_lock()***

© 2004, QNX Software Systems Ltd.

*Lock a resource*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_lock(resmgr_context_t * ctp,
 io_lock_t * msg,
 iofunc_ocb_t * ocb,
 iofunc_attr_t * attr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_lock_t` structure that contains the message that the resource manager received.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The function `iofunc_lock()` does what is required for POSIX locks. This function isn't currently implemented.

### **Returns:**

`ENOSYS` The `iofunc_lock()` function isn't currently supported.



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_lock\_alloc()*, *iofunc\_lock\_free()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_lock\_alloc()***

© 2004, QNX Software Systems Ltd.

*Allocate memory to lock structures*

### **Synopsis:**

```
#include <sys/iofunc.h>

iofunc_lock_list_t *iofunc_lock_alloc
(resmgr_context_t *ctp,
 IOFUNC_OCB_T *ocb,
 size_t size);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*ocb* A pointer to the the Open Control Block (typically a `iofunc_ocb_t` structure) that was created when the client opened the resource.

*size* The amount of memory that you want to allocate.

### **Library:**

`libc`

### **Description:**

The function `iofunc_lock_alloc()` is used by `iofunc_lock()` to allocate memory to lock structures.

### **Returns:**

A pointer to a zeroed buffer that the POSIX layer uses for locks, or NULL if no memory could be allocated.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_lock()*, *iofunc\_lock\_free()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_lock\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_LOCK messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_lock_default(resmgr_context_t * ctp,
 io_lock_t * msg,
 iofunc_ocb_t * ocb);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_lock_t` structure that contains the message that the resource manager received; see `iofunc_lock()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_lock_default()` function implements POSIX semantics for the `_IO_LOCK` message (generated as a result of a client `fcntl()` call).

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `lock` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_lock_default()` function verifies that the client has the necessary permissions to effect a lock on the resource. This includes checking for read and write permissions against the type of lock being effected. If so, the lock is performed, modifying elements of the `ocb->attr` structure, and updating `ocb->attr->locklist` to reflect the

new lock. This function calls *iofunc\_lock()*, which does the actual work.



---

If your resource manager calls *iofunc\_lock\_default()*, it must call *iofunc\_close\_dup\_default()* and *iofunc\_unblock\_default()* in their respective handlers.

---

## Returns:

|        |                                                                                                                    |
|--------|--------------------------------------------------------------------------------------------------------------------|
| EOK    | Successful completion.                                                                                             |
| EINVAL | An invalid range was specified for the lock operation, or an invalid lock operation was attempted.                 |
| EBADF  | An attempt to perform a read lock on a write-only resource, or a write lock on a read-only resource was attempted. |
| ENOMEM | Insufficient memory exists to allocate an internal lock structure.                                                 |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_func\_init()*, *iofunc\_lock()*, **iofunc\_ocb\_t**,  
*iofunc\_time\_update()*, *resmgr\_attach()*, **resmgr\_context\_t**,  
**resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>

void iofunc_lock_free(iofunc_lock_list_t* lock,
 size_t size);
```

### **Arguments:**

*lock*     A pointer to the `iofunc_lock_list_t` list that you want to free.

*size*     The amount of memory that you want to free.

### **Library:**

`libc`

### **Description:**

The function `iofunc_lock_free()` frees `lock` structures allocated by `iofunc_lock_alloc()`.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_lock()*, *iofunc\_lock\_calloc()*

Writing a Resource Manager chapter of the *Programmer's Guide*.



## ***iofunc\_lock\_ocr\_default()***

*Default handler for the lock\_ocr callout*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_lock_ocr_default(resmgr_context_t *ctp,
 void *reserved,
 iofunc_ocr_t *ocr);
```

### **Arguments:**

*ctp*            A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*reserved*      This argument must be NULL.

*ocr*            A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_lock_ocr_default()` function calls `iofunc_attr_lock()` to enforce locking on the attributes for the group of messages that were sent by the client.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `lock_ocr` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

### **Returns:**

EOK            Success.

EAGAIN        On the first use, all kernel mutex objects were in use.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_lock()*, *iofunc\_func\_init()*, *iofunc\_ocb\_t*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_lseek (resmgr_context_t* ctp,
 io_lseek_t* msg,
 iofunc_ocb_t* ocb,
 iofunc_attr_t* attr);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_lseek_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

## Library:

`libc`

## Description:

The `iofunc_lseek()` helper function implements POSIX semantics for the client's `lseek()` call, which is received as an `_IO_LSEEK` message by the resource manager.

The `iofunc_lseek()` function handles the three different *whence* cases: `SEEK_SET`, `SEEK_CUR`, and `SEEK_END`, updating the `ocb->offset` field with the new position.

Note that if the IOFUNC\_MOUNT\_32BIT flag isn't set in the mount structure, *iofunc\_lseek()* handles 64-bit position offsets. If the flag is set (meaning this device supports only 32-bit offsets), the resulting offset value is treated as a 32-bit offset, and if it overflows 32 bits, it's truncated to LONG\_MAX. Also, this function handles combine messages correctly, simplifying the work required to support lseek.

### **io\_lseek\_t structure**

The `io_lseek_t` structure holds the `_IO_LSEEK` message received by the resource manager:

```
struct _io_lseek {
 uint16_t type;
 uint16_t combine_len;
 short whence;
 uint16_t zero;
 uint64_t offset;
};

typedef union {
 struct _io_lseek i;
 uint64_t o;
} io_lseek_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The *i* member is a structure of type `_io_lseek` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_LSEEK</code> .                                                                                                                                                                                            |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |
| <i>whence</i>      | <code>SEEK_SET</code> , <code>SEEK_CUR</code> , or <code>SEEK_END</code> .                                                                                                                                          |

*offset*            The relative offset from the file position determined by the *whence* member.

The *o* member is the offset after the operation is complete.

## Returns:

EOK            Successful completion.

EINVAL        The *whence* member in the `_IO_LSEEK` message wasn't one of `SEEK_SET`, `SEEK_CUR`, or `SEEK_END`, or the resulting position after the offset was applied resulted in a negative number (overflow).

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

`iofunc_attr_t`, `iofunc_lseek_default()`, `iofunc_ocb_t`, `lseek()`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_lseek\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_LSEEK messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_lseek_default(resmgr_context_t* ctp,
 io_lseek_t* msg,
 iofunc_ocr_t* ocb);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_lseek_t` structure that contains the message that the resource manager received. For more information, see the documentation for *iofunc\_lseek()*.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The *iofunc\_lseek\_default()* function implements POSIX semantics for the client's *lseek()* call, which is received as an `_IO_LSEEK` message by the resource manager.

You can place this function directly into the *io\_funcs* table passed to *resmgr\_attach()*, at the *lseek* position, or you can call *iofunc\_func\_init()* to initialize all of the functions to their default values.

The *iofunc\_lseek\_default()* function calls *iofunc\_lseek()* to do the actual work, and (if installed in the *io\_funcs* table) issues the reply back to the client.

## Returns:

- |        |                                                                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EOK    | Successful completion.                                                                                                                                                                                                                                   |
| EINVAL | The <i>whence</i> member in the <code>_IO_LSEEK</code> message wasn't one of <code>SEEK_SET</code> , <code>SEEK_CUR</code> , or <code>SEEK_END</code> , or the resulting position after the offset was applied resulted in a negative number (overflow). |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_lseek()*, `iofunc_ocb_t`, *lseek()*, *resmgr\_attach()*, `resmgr_context_t`, `resmgr_io_funcs_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_mknod()***

© 2004, QNX Software Systems Ltd.

*Verify a client's ability to make a new filesystem entry point*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_mknod(resmgr_context_t *ctp,
 io_mknod_t *msg,
 iofunc_attr_t *attr,
 iofunc_attr_t *datr,
 struct _client_info *info);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_mknod_t` structure that contains the message that the resource manager received; see below.
- attr* NULL, or a pointer to the `iofunc_attr_t` structure that describes the characteristics of the resource.
- datr* A pointer to the `iofunc_attr_t` structure that you must set. The `iofunc_attr_t` structure describes the attributes of the parent directory.
- info* NULL, or a pointer to a `_client_info` structure that contains information about the client. For information about this structure, see *ConnectClientInfo()*.

### **Library:**

`libc`

### **Description:**

The *iofunc\_mknod()* helper function supports *mknod()* requests by verifying that the client can make a new filesystem entry point. It's similar to *iofunc\_open()*.



The *iofunc\_mknod()* function checks to see if the client (described by the optional *info* structure) has access to open the resource (name passed in the *msg* structure). The *attr* structure describes the resource's attributes, and the optional *dattr* structure defines the attributes of the parent directory (i.e. if *dattr* isn't NULL, it implies that the resource identified by *attr* is being created within the directory specified by *dattr*).

The *info* argument can be passed as NULL, in which case *iofunc\_mknod()* obtains the client information itself via a call to *iofunc\_client\_info()*. It is, of course, more efficient to get the client info once, rather than calling this function with NULL every time.

If an error occurs, the function returns information about a client's connection in *info* and a constant.

### **io\_mknod\_t structure**

The *io\_mknod\_t* structure holds the *\_IO\_CONNECT* message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_ftype_reply ftype_reply;
} io_mknod_t;
```

This message structure is a union of an input message (coming to the resource manager), *\_io\_connect*, and two possible output or reply messages (going back to the client):

- *\_io\_connect\_link\_reply* if the reply is redirecting the client to another resource

Or:

- *\_io\_connect\_ftype\_reply* if the reply consists of a status and a file type.

**Returns:**

|          |                                                                          |
|----------|--------------------------------------------------------------------------|
| EOK      | Success.                                                                 |
| EBADFSYS | NULL was passed in <i>dattr</i> .                                        |
| EFAULT   | A fault occurred when the kernel tried to access the <i>info</i> buffer. |
| EINVAL   | The client process is no longer valid.                                   |
| ENOSYS   | NULL was passed in <i>info</i> .                                         |
| EPERM    | The group ID or owner ID didn't match.                                   |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`_io_connect`, `_io_connect_link_reply`,  
`_io_connect_fstype_reply`, `iofunc_client_info()`, `iofunc_open()`,  
`mknod()`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_mmap (resmgr_context_t * hdr,
 io_mmap_t * msg,
 iofunc_ocb_t * ocb,
 iofunc_attr_t * attr);
```

## Arguments:

- hdr* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_mmap_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

## Library:

`libc`

## Description:

The `iofunc_mmap()` helper function provides functionality for the `_IO_MMAP` message. The `_IO_MMAP` message is an outcall from the Memory Manager (a part of the QNX Neutrino microkernel's `procnto`).

Note that if the Process Manager is to be able to execute from this resource, then you must use the `iofunc_mmap()` function.

**io\_mmap\_t structure**

The `io_mmap_t` structure holds the `_IO_MMAP` message received by the resource manager:

```
struct _io_mmap {
 uint16_t type;
 uint16_t combine_len;
 uint32_t prot;
 uint64_t offset;
 struct _msg_info info;
 uint32_t zero[6];
};

struct _io_mmap_reply {
 uint32_t zero;
 uint32_t flags;
 uint64_t offset;
 int32_t coid;
 int32_t fd;
};

typedef union {
 struct _io_mmap i;
 struct _io_mmap_reply o;
} io_mmap_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The *i* member is a structure of type `_io_mmap` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_MMAP</code> .                                                                                                                                                                                             |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> . |
| <i>prot</i>        | The access capabilities that the client wants to use for the memory region being mapped. This can be a                                                                                                              |

combination of at least the following protection bits, as defined in `<sys/mman.h>`:

- `PROT_EXEC` — the region can be executed.
- `PROT_NOCACHE` — disable caching of the region (e.g. so it can be used to access dual-ported memory).
- `PROT_NONE` — the region can't be accessed.
- `PROT_READ` — the region can be read.
- `PROT_WRITE` — the region can be written.

*offset*      The offset into shared memory of the location that the client wants to start mapping.

*info*        A pointer to a `_msg_info`, structure that contains information about the message received by the resource manager.

The *o* member of the `io_mmap_t` structure is a structure of type `_io_mmap_reply` that contains the following members:

*flags*      Reserved for future use.

*offset*     Reserved for future use.

*coid*       A file descriptor that the process manager can use to access the mapped file.

*fd*         Reserved for future use.

## Returns:

A nonpositive value (i.e.  $\leq 0$ )

Successful completion.

EROFS      An attempt to memory map (mmap) a read-only file, using the `PROT_WRITE` page protection mode.

EACCES     The client doesn't have the appropriate permissions.

ENOMEM      Insufficient memory exists to allocate internal resources required to effect the mapping.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_attr_t`, `iofunc_mmap_default()`, `iofunc_ocb_t`,  
`_msg_info`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_mmap_default (resmgr_context_t * hdr,
 io_mmap_t * msg,
 iofunc_ocb_t * ocb);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_mmap_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_mmap()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

## Library:

`libc`

## Description:

The `iofunc_mmap_default()` function provides functionality for the `_IO_MMAP` message. This message is private to the Memory Manager (a part of the Neutrino microkernel's `procnto`).

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `mmap` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

Note that if the Process Manager is to be able to execute from this resource, then you must use the `iofunc_mmap()` function.

The *iofunc\_mmap\_default()* function calls *iofunc\_mmap()*, to do the actual work, and (if installed in the *io\_funcs* table) issues the reply back to the client.

## Returns:

A nonpositive value (i.e.  $\leq 0$ )

Successful completion.

EROFS      An attempt to memory map (mmap) a read-only file, using the PROT\_WRITE page protection mode.

EACCES     The client doesn't have the appropriate permissions.

ENOMEM    Insufficient memory exists to allocate internal resources required to effect the mapping.

## Classification:

QNX Neutrino

### Safety

---

Cancellation point    No

Interrupt handler     No

Signal handler        Yes

Thread                 Yes

## See also:

*iofunc\_func\_init()*, *iofunc\_mmap()*, *iofunc\_ocb\_t*, *resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.



**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_notify(resmgr_context_t *ctp,
 io_notify_t *msg,
 iofunc_notify_t *nop,
 int trig,
 const int *notifycounts,
 int *armed);
```

**Arguments:**

- |             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>  | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>msg</i>  | A pointer to the <code>io_notify_t</code> structure that contains the message that the resource manager received; see below.                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>nop</i>  | <p>An array of <code>iofunc_notify_t</code> structures that represent the events supported by the calling resource manager. Traditionally this array contained three members which represent, in order, the input, output, and out-of-band notification lists. Since the addition of extended events (see below), three is now the minimum size of this array. The actual size must support indexing by the conditions being triggered up to <code>_NOTIFY_MAXCOND</code>.</p> <p>Generally, this structure is maintained by the resource manager within an extended attributes structure.</p> |
| <i>trig</i> | A bitmask indicating which sources are currently satisfied, and could cause a trigger to occur. This bitmask may be indicated via two sets of flags. Traditionally, the value was any combination of <code>_NOTIFY_COND_INPUT</code> , <code>_NOTIFY_COND_OUTPUT</code>                                                                                                                                                                                                                                                                                                                        |

and `_NOTIFY_COND_OBAND`. With the addition of extended events, this can also be any combination of the `_NOTIFY_CONDE*` flags. Note the following flags are considered equivalent:

```
_NOTIFY_COND_INPUT == _NOTIFY_CONDE_RDNORM
_NOTIFY_COND_OUTPUT == _NOTIFY_CONDE_WRNORM
_NOTIFY_COND_OBAND == _NOTIFY_CONDE_RDBAND
```

Setting the `_NOTIFY_COND_EXTEN` flag may affect the “armed” parameter as indicated below.

You typically set this value, based on the conditions in effect at the time of the call.

*notifycounts* NULL, or an array of integers representing the number of elements that must be present in the queue of each event represented by the *nop* array in order for the event to be triggered. Both this array and the *nop* array should contain the same number of elements. Note that if any condition is met, nothing is armed. Only if none of the conditions are met, does the event get armed in accordance with the *notifycounts* parameter. If this parameter isn't specified (passed as NULL), a value of 1 is assumed for all counts.

*armed* NULL, or a pointer to a location where the function can store a 1 to indicate that a notification entry is armed, or a 0 otherwise. If the `_NOTIFY_COND_EXTEN` bit is set in the *trig* parameter and *armed* is not NULL, it is promoted from being a strictly resultant parameter to value resultant and must contain the number of elements in the *nop* and *notifycounts* array (provided *notifycounts* is not NULL) at the time of the call. If either of the above conditions is not met, *armed* remains strictly a resultant parameter, and the traditional number of three elements is assumed in *nop* and *notifycounts*.

## Library:

`libc`

## Description:

The POSIX layer helper function *iofunc\_notify()* is used by a resource manager to implement notification.

This routine examines the message that the resource manager received (passed in the *msg* argument), and determines what action the client code is attempting to perform:

### `_NOTIFY_ACTION_POLL`

Return a one-part IOV with the *flags* field set to indicate which conditions (input, output, or out-of-band) are available. The caller should return (*\_RESMGR\_NPARTS(1)*) to the resource manager library, which returns a one-part message to the client.

### `_NOTIFY_ACTION_POLLARM`

Similar to `_NOTIFY_ACTION_POLL`, with the additional characteristic of arming the event if *none* of the conditions is met.

### `_NOTIFY_ACTION_TRANARM`

For each of the sources specified, create a notification entry and store the client's `struct sigevent` event structure in it. Note that only one transition arm is allowed at a time *per device*. If the client specifies an event of `SIGEV_NONE`, the action is to disarm. When the event is triggered, the notification is automatically disarmed.

## `io_notify_t` structure

The `io_notify_t` structure holds the `_IO_NOTIFY` message received by the resource manager:

```
struct _io_notify {
 uint16_t
 uint16_t
 type;
 combine_len;
```

```
 int32_t action;
 int32_t flags;
 struct sigevent event;
};

struct _io_notify_reply {
 uint32_t zero;
 uint32_t flags;
};

typedef union {
 struct _io_notify i;
 struct _io_notify_reply o;
} io_notify_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The *i* member is a structure of type `_io_notify` that contains the following members:

|                    |                                                                                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_NOTIFY</code> .                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> .                                                                                                                                                                                       |
| <i>action</i>      | <code>_NOTIFY_ACTION_POLL</code> , <code>_NOTIFY_ACTION_POLLARM</code> , or <code>_NOTIFY_ACTION_TRANARM</code> , as described above.                                                                                                                                                                                                                                                                     |
| <i>flags</i>       | One of the following: <ul style="list-style-type: none"><li>• <code>_NOTIFY_COND_INPUT</code> — this condition is met when there are one or more units of input data available (i.e. clients can now issue reads).</li><li>• <code>_NOTIFY_COND_OUTPUT</code> — this condition is met when there’s room in the output buffer for one or more units of data (i.e. clients can now issue writes).</li></ul> |

- `_NOTIFY_COND_OBAND` — the condition is met when one or more units of out-of-band data are available.

*event* A pointer to a **sigevent** structure that defines the event that the resource manager is to deliver once a condition is met.

The *o* member is a structure of type `_io_notify_reply` that contains the following members:

*flags* Which of the conditions were triggered; see the *flags* for `_io_notify`, above.

### **iofunc\_notify\_t structure**

The `iofunc_notify_t` structure is defined in `<sys/iofunc.h>` as follows:

```
typedef struct _iofunc_notify {
 int cnt;
 struct _iofunc_notify_event *list;
} iofunc_notify_t;
```

The members of the `iofunc_notify_t` structure include:

*cnt* The smallest *cnt* member in the list; see below.

*list* A pointer to a linked list of `iofunc_notify_event_t` structures that represent (in order), the input, output, and out-of-band notification lists.

The `iofunc_notify_event_t` is defined as:

```
typedef struct _iofunc_notify_event {
 struct _iofunc_notify_event *next;
 int rvid;
 int scoid;
 int cnt;
 struct sigevent event;
} iofunc_notify_event_t;
```

The members of the `iofunc_notify_event_t` structure include:

- next*      A pointer to the next element in the list.
- rcvid*     The receive ID of the client to notify.
- scoid*     The server connection ID.
- cnt*       The number of bytes available. Some clients, such as `io-char`, may want a sufficiently large amount of data to be available before they access it.
- event*     A pointer to a `sigevent` structure that defines the event that the resource manager is to deliver once a condition is met.

**Returns:**

`EBUSY`      A notification was already armed for this resource, and this library function enforces a restriction of one per resource.

`_RESMGR_NPARTS (1)`

Normal return, indicates a one-part IOV should be returned to the client.

**Examples:**

See the Writing a Resource Manager chapter of *Programmer's Guide*.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*iofunc\_notify\_remove()*, *iofunc\_notify\_trigger()*, *\_RESMGR\_NPARTS()*,  
**sigevent**

“Handling *ionotify()* and *select()*” in the Writing a Resource Manager  
chapter of the *Programmer’s Guide*.

## ***iofunc\_notify\_remove()***

© 2004, QNX Software Systems Ltd.

*Remove notification entries from list*

### **Synopsis:**

```
#include <sys/iofunc.h>

void iofunc_notify_remove(resmgr_context_t * ctp,
 iofunc_notify_t * nop);
```

### **Arguments:**

*ctp* NULL, or a pointer to a **resmgr\_context\_t** structure for the client whose entries you want to remove.

*nop* An array of three **iofunc\_notify\_t** structures that represent (in order), the input, output, and out-of-band notification lists whose entries you want to remove; for information about this structure, see the documentation for *iofunc\_notify()*.

### **Library:**

**libc**

### **Description:**

The *iofunc\_notify\_remove()* function removes all of the entries associated with the current client from the notification list passed in *nop*. The client information is obtained from the *ctp*.

If the *ctp* pointer is NULL, then *all* of the notify entries will be removed. A resource manager generally calls this function, with NULL as the *ctp* in the *close\_ocb* callout, to clean up all handles associated with this connection. If the handles are shared between several connections, then the *ctp* should be provided to clean up after each client.

### **Examples:**

See the “Writing a Resource Manager” chapter in the *Programmer’s Guide*.



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_notify()*, *iofunc\_notify\_trigger()*

“Handling *ionotify()* and *select()*” in the Writing a Resource Manager chapter of the *Programmer’s Guide*.

## ***iofunc\_notify\_trigger()***

© 2004, QNX Software Systems Ltd.

*Send notifications to queued clients*

### **Synopsis:**

```
#include <sys/iofunc.h>

void iofunc_notify_trigger(iofunc_notify_t *nop,
 int count,
 int index);
```

### **Arguments:**

- nop* An array of three `iofunc_notify_t` structures that represent (in order), the input, output, and out-of-band notification lists whose entries you want to examine; for information about this structure, see the documentation for `iofunc_notify()`.
- count* The count that you want to compare to the trigger value for the event.
- index* The index into the *nop* array that you want to check; one of the following:
- IOFUNC\_NOTIFY\_INPUT
  - IOFUNC\_NOTIFY\_OUTPUT
  - IOFUNC\_NOTIFY\_OBAND

### **Library:**

`libc`

### **Description:**

The `iofunc_notify_trigger()` function examines all entries given in the list maintained at `nop [index]` to see if the event should be delivered to the client. If the specified `count` is greater than the trigger count for the particular notification list element, this function calls `MsgDeliverEvent()` to deliver the event to the client whose `rcvid` is stored in the notification list element, and the list element is disarmed.

Note that if the client has specified a code of `SI_NOTIFY`, then the value that the client specified (e.g. the *value* member of the `struct sigevent`) has the top three bits ORed with the reason for the trigger (this is the expression `_NOTIFY_COND_INPUT << index`), as in the following table:

```
index = IOFUNC_NOTIFY_INPUT
 0x10000000, or _NOTIFY_COND_INPUT

index = IOFUNC_NOTIFY_OUTPUT
 0x20000000, or _NOTIFY_COND_OUTPUT

index = IOFUNC_NOTIFY_OBAND
 0x40000000, or _NOTIFY_COND_OBAND
```

If the client has specified a code of something other than `SI_NOTIFY` then this routine doesn't modify the value member in any way.

## Examples:

See the Writing a Resource Manager chapter of *Programmer's Guide*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_notify()*, *iofunc\_notify\_remove()*, **sigevent**

“Handling *ionotify()* and *select()*” in the Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_ocb_attach(
 resmgr_context_t * ctp,
 io_open_t * msg,
 iofunc_ocb_t * ocb,
 iofunc_attr_t * attr,
 const resmgr_io_funcs_t * io_funcs);
```

**Arguments:**

|                 |                                                                                                                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions.                                                 |
| <i>msg</i>      | A pointer to the <code>io_open_t</code> structure that contains the message that the resource manager received. For more information, see the documentation for <code>iofunc_open()</code> . |
| <i>ocb</i>      | NULL, or a pointer to the <code>iofunc_ocb_t</code> structure for the Open Control Block that was created when the client opened the resource.                                               |
| <i>attr</i>     | A pointer to a <code>iofunc_attr_t</code> structure that defines the characteristics of the device that the resource manager handles.                                                        |
| <i>io_funcs</i> | A pointer to a <code>resmgr_io_funcs_t</code> that specifies the I/O functions for the resource manager.                                                                                     |

**Library:**

`libc`

## Description:

The *iofunc\_ocb\_attach()* function examines the mode specified by the *io\_open msg*, and increments the read and write count flags (*ocb->attr->rcount* and *ocb->attr->>wcount*), and the locking flags (*ocb->attr->rlocks* and *ocb->attr->>wlocks*), as specified by the open mode.

This function is called by *iofunc\_open\_default()* as part of its initialization.

This function allocates the memory for the OCB if you pass NULL as the *ocb*.

## Returns:

EOK      Successful completion.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_attr\_init()*, *iofunc\_attr\_t*, *iofunc\_ocb\_detach()*,  
*iofunc\_ocb\_t*, *iofunc\_open\_default()*, *resmgr\_context\_t*,  
*resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

iofunc_ocb_t * iofunc_ocb_alloc(
 resmgr_context_t * ctp,
 iofunc_attr_t * attr);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- attr* A pointer to a `iofunc_attr_t` structure that defines the characteristics of the device that the resource manager handles.

## Library:

`libc`

## Description:

The `iofunc_ocb_alloc()` function allocates an iofunc OCB. It has a number of uses:

- It can be used as a helper function to encapsulate the allocation of the iofunc OCB, so that your routines don't have to know the details of the iofunc OCB structure.
- Because it's in the resource manager shared library, you can override this function with your own, allowing you to manage an OCB that has additional members, perhaps specific to your particular resource manager. If you do this, be sure to place the iofunc OCB structure as the first element of your extended OCB, and also override the `iofunc_ocb_free()` function to release memory.
- Another reason to override `iofunc_ocb_alloc()` might be to place limits on the number of OCBs that are in existence at any one

time; the current function simply allocates OCBs until the free store is exhausted.



---

You should fill in the attribute's mount structure (i.e. the *attr->mount* pointer) instead of replacing this function.

If you specify *iofunc\_ocb\_alloc()* and *iofunc\_ocb\_free()* callouts in the attribute's mount structure, then you should use the callouts instead of calling the standard *iofunc\_ocb\_alloc()* and *iofunc\_ocb\_free()* functions.

---

## Returns:

A pointer to an `iofunc_ocb_t` OCB structure.

## Examples:

Override *iofunc\_ocb\_alloc()* and *iofunc\_ocb\_free()* to manage an extended OCB:

```
typedef struct
{
 iofunc_ocb_t iofuncOCB; /* the OCB used by iofunc_* */
 int myFlags;
 char moreOfMyStuff;
} MyOCBT;

MyOCBT *iofunc_ocb_alloc (resmgr_context_t *ctp,
 iofunc_attr_t *attr)
{
 return ((MyOCBT *) calloc (1, sizeof (MyOCBT)));
}

void iofunc_ocb_free (MyOCBT *ocb)
{
 free (ocb);
}
```



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_ocb\_free()*, `iofunc_ocb_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_ocb\_detach()***

© 2004, QNX Software Systems Ltd.

*Release Open Control Block resources*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_ocb_detach(resmgr_context_t * ctp,
 iofunc_ocb_t * ocb);
```

### **Arguments:**

*ctp* A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions.

*ocb* A pointer to the **iofunc\_ocb\_t** structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

**libc**

### **Description:**

The *iofunc\_ocb\_detach()* function releases any resources allocated to the passed *ocb*, such as any memory map (mmap) entries.



---

This function doesn't free the memory associated with the OCB itself.

---

The *iofunc\_ocb\_detach()* function also updates the time structure, by calling *iofunc\_time\_update()*, and decrements the read, write, lock, and use counters, according to the mode that was used to open the resource (from *ocb->ioflag*).

The counters are incremented in *iofunc\_ocb\_attach()*, and represent the number of OCBs that are using the managed resource in the respective manners (e.g.: *ocb->attr->rcount* keeps count of how many OCBs are using the resource specified by *attr* for read access).

If you're are using *iofunc\_mmap()* or *iofunc\_mmap\_default()*, you must call *iofunc\_ocb\_detach()* to clean up. This function is called by *iofunc\_close\_ocb()*.

## Returns:

A bitwise OR of flags describing the state of the managed resource:

### IOFUNC\_OCB\_LAST\_READER

This OCB was the last one performing read operations on the resource. This flag is set when the *ocb->attr->rcount* flag is decremented to zero.

### IOFUNC\_OCB\_LAST\_WRITER

This OCB was the last one performing write operations on the resource. This flag is set when the *ocb->attr->>wcount* flag is decremented to zero.

### IOFUNC\_OCB\_LAST\_RDLOCK

This OCB was the last one holding a read lock on the resource. This flag is set when the *ocb->attr->rlocks* flag is decremented to zero.

### IOFUNC\_OCB\_LAST\_WRLOCK

This OCB was the last one holding a write lock on the resource. This flag is set when the *ocb->attr->>wlocks* flag is decremented to zero.

### IOFUNC\_OCB\_LAST\_INUSE

This OCB was the last one using the resource. This flag is set when the *ocb->attr->count* flag is decremented to zero.

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_close\_ocb()*, *iofunc\_close\_ocb\_default()*, *iofunc\_mmap()*,  
*iofunc\_mmap\_default()*, *iofunc\_ocb\_attach()*, **iofunc\_ocb\_t**,  
*iofunc\_time\_update()*, **resmgr\_context\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>

void iofunc_ocb_free(iofunc_ocb_t * ocb);
```

### **Arguments:**

*ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_ocb_free()` function returns the memory allocated to an iofunc OCB to the free store pool. This function is the complement of `iofunc_ocb_calloc()`.

If you've overridden the definition of `iofunc_ocb_calloc()`, you should also override the definition of `iofunc_ocb_free()` to correctly handle the release of the memory. This is because the `iofunc_ocb_calloc()` functions uses an internal memory management function to allocate the memory, and the default `iofunc_ocb_free()` function also uses this internal function to deallocate memory. Therefore, you can't mix internal memory management functions (`_scalloc()` and `_sfree()`) with user-level memory management functions (`calloc()` and `free()`).



---

You should fill in the attribute's mount structure (i.e. the *attr->mount* pointer) instead of replacing this function.

If you specify *iofunc\_ocb\_free()* and *iofunc\_ocb\_alloc()* callouts in the attribute's mount structure, then you should use the callouts instead of calling the standard *iofunc\_ocb\_free()* and *iofunc\_ocb\_alloc()* functions.

---

## Examples:

See *iofunc\_ocb\_alloc()*.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_close\_ocb()*, *iofunc\_ocb\_alloc()*, *iofunc\_ocb\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```

#include <sys/iofunc.h>

typedef struct _iofunc_ocb {
 IOFUNC_ATTR_T *attr; /* Used to find iofunc_ocb */
 int32_t ioflag; /* open's oflag + 1 */
#if !defined(_IOFUNC_OFFSET_BITS) || _IOFUNC_OFFSET_BITS == 64
 #if _FILE_OFFSET_BITS - 0 == 64
 off_t offset;
 #else
 off64_t offset;
 #endif
#elif _IOFUNC_OFFSET_BITS - 0 == 32
 #if !defined(_FILE_OFFSET_BITS) || _FILE_OFFSET_BITS == 32
 #if defined(__LITTLEENDIAN__)
 off_t offset;
 off_t offset_hi;
 #elif defined(__BIGENDIAN__)
 off_t offset_hi;
 off_t offset;
 #else
 #error endian not configured for system
 #endif
 #else
 #if defined(__LITTLEENDIAN__)
 int32_t offset;
 int32_t offset_hi;
 #elif defined(__BIGENDIAN__)
 int32_t offset_hi;
 int32_t offset;
 #else
 #error endian not configured for system
 #endif
 #endif
 #endif
 uint16_t sflag;
 uint16_t flags;
 void *reserved;
} iofunc_ocb_t;

```

## Description:

The `iofunc_ocb_t` structure is an *Open Control Block*, a block of data that's established by a resource manager during its handling of the client's `open()` function.

A resource manager creates an instance of this structure whenever a client opens a resource. For example, `iofunc_open_default()` calls `iofunc_ocb_calloc()` to allocate an OCB. The OCB exists until the client closes the file descriptor associated with the open operation. The resource manager passes this structure to all of the functions that implement the I/O operations for the file descriptor.

The `iofunc_ocb_t` structure includes the following members:

*attr* A pointer to the OCB's attributes. By default, this structure is of type `iofunc_attr_t`, but you can redefine the `IOFUNC_ATTR_T` manifest if you want to use a different structure in your resource manager.

*ioflag* The mode (e.g. reading, writing, blocking) that the resource was opened with.



---

The bits in this member are the same as those for the *ioflag* argument to `open()` plus 1.

---

This information is inherited from the `io_connect_t` structure that's available in the message passed to the open handler.

*offset, offset\_hi*

The read/write offset into the resource (e.g. our current `lseek()` position within a file), defined in a variety of ways to suit 32- and 64-bit offsets. Your resource manager can modify this offset.

*sflag* The sharing mode; see `sopen()`. This information is inherited from the `io_connect_t` structure that's available in the message passed to the open handler.



*flags* When the IOFUNC\_OCB\_PRIVILEGED bit is set, a privileged process (i.e. `root`) performed the `open()`. Additionally, you can use flags in the range defined by IOFUNC\_OCB\_FLAGS\_PRIVATE (see `<sys/iofunc.h>`) for your own purposes. Your resource manager can modify these flags.

### **Classification:**

QNX Neutrino

### **See also:**

`iofunc_attr_t`, `iofunc_ocb_calloc()`, `iofunc_open_default()`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_open()***

© 2004, QNX Software Systems Ltd.

*Verify a client's ability to open a resource*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_open(resmgr_context_t *ctp,
 io_open_t *msg,
 iofunc_attr_t *attr,
 iofunc_attr_t *datr,
 struct _client_info *info);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_open_t` structure that contains the message that the resource manager received; see below.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the resource.
- datr* NULL, or a pointer to the `iofunc_attr_t` structure that describes the characteristics of the parent directory.
- info* NULL, or a pointer to a `_client_info` structure that contains the information about a client connection. For information about this structure, see *ConnectClientInfo()*.

### **Library:**

`libc`

### **Description:**

The *iofunc\_open()* function checks to see if the client (described by the optional *info* structure) has access to open the resource whose name is passed in *msg->connect.path*.

The *attr* structure describes the resource's attributes. The optional *datr* structure defines the attributes of the parent directory; if *datr*

isn't NULL, the resource identified by *attr* is being created within the directory specified by *dattr*.

The *info* argument can be passed as NULL, in which case *iofunc\_open()* obtains the client information itself via a call to *iofunc\_client\_info()*. It is, of course, more efficient to get the client info once, rather than calling this function with NULL every time.

Note that if you're handling a request to read directory entry, you must return data formatted to match the **struct dirent** type. A helper function, *iofunc\_stat()*, can aid in this.

A resource manager's response to an *open()* request isn't always a yes-or-no answer. It's possible to return a connect message indicating that the server would like some other action taken. For example, if the open occurs on a path that represents a symbolic link to some other path, the server could respond using the *\_IO\_SET\_CONNECT\_RET()* macro and the *\_IO\_CONNECT\_RET\_LINK* value.

For example, an open handler that only redirects pathnames might look something like:

```
io_open(resmgr_context_t *ctp, io_open_t *msg,
 iofunc_attr_t *dattr, void *extra) {
 char *newpath;

 /* Do all the error/access checking ... */

 /* Lookup the redirected path and store
 the new path in 'newpath' */
 newpath = get_a_new_path(msg->connect.path);

 _IO_SET_CONNECT_RET(ctp, _IO_CONNECT_RET_LINK);
 len = strlen(newpath) + 1;

 msg->link_reply.eflag = msg->connect.eflag;
 msg->link_reply.nentries = 0;
 msg->link_reply.path_len = len;
 strcpy((char *) (msg->link_reply + 1), newpath);

 len += sizeof(msg->link_reply);

 return(_RESMGR_PTR(ctp, &msg->link_reply, len));
}
```

In this example, we use the macro `_IO_SET_CONNECT_RET()` (defined in `<sys/iomsg.h>`) to set the `ctp->status` field to `_IO_CONNECT_RET_LINK`. This value indicates to the resource-manager framework that the return value isn't actually a simple return code, but a new request to be processed.

The path for this new request follows directly after the `link_reply` structure and is `path_len` bytes long. The final few lines of the code just stuff an IOV with the reply message (and the new path to be queried) and return to the resource-manager framework.

## **io\_open\_t structure**

The `io_open_t` structure holds the `_IO_CONNECT` message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_fstype_reply fstype_reply;
} io_open_t;
```

This message structure is a union of an input message (coming to the resource manager), `_io_connect`, and two possible output or reply messages (going back to the client):

- `_io_connect_link_reply` if the reply is redirecting the client to another resource

Or:

- `_io_connect_fstype_reply` if the reply consists of a status and a file type.

## **Returns:**

|       |                                                                                                                                  |
|-------|----------------------------------------------------------------------------------------------------------------------------------|
| EOK   | Successful completion.                                                                                                           |
| Other | There was an error, as defined by the POSIX semantics for the open call. This error should be returned to the next higher level. |

**Examples:**

This is a sample skeleton for a typical filesystem, in pseudo-code, to illustrate the steps that need to be taken to handle an open request for a file:

```

if the open request is for a path (i.e. multiple
directory levels)
 call iofunc_client_info to get information
 about client
 for each directory component
 call iofunc_check_access to check execute
 permission for access
 /*
 recall that execute permission on a
 directory is really the "search"
 permission for that directory
 */
 next
/*
at this point you have verified access
to the target
*/
endif

if O_CREAT is set and the file doesn't exist
 call iofunc_open, passing the attribute of the
 parent as dattr
 if the iofunc_open succeeds,
 do the work to create the new inode,
 or whatever
 endif
else
 call iofunc_open, passing the attr of the file
 and NULL for dattr
endif

/*
at this point, check for things like o_trunc,
etc. -- things that you have to do for the attr
*/

call iofunc_ocb_attach
return EOK

```

For a device (i.e. *resmgr\_attach()* didn't specify that the managed resource is a directory), the following steps apply:

```
/*
```

```
 at startup time (i.e.: in the main() of the
 resource manager)
*/
call iofunc_attr_init to initialize an attribute
 structure

/* in the io_open message handler: */
call iofunc_open, passing in the attribute of the
 device and NULL for dattr

call iofunc_ocb_attach
return EOK
```

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

`_io_connect`, `_io_connect_link_reply`,  
`_io_connect_fstype_reply`, `iofunc_attr_init()`,  
`iofunc_check_access()`, `iofunc_client_info()`, `iofunc_ocb_attach()`,  
`iofunc_stat()`, `resmgr_open_bind()`

Writing a Resource Manager chapter of the *Programmer's Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_open_default(resmgr_context_t *ctp,
 io_open_t *msg,
 iofunc_attr_t *attr,
 void *extra);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_open_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_open()`.
- attr* A pointer to the `iofunc_attr_t` structure that defines the characteristics of the device that the resource manager is controlling.

### **Library:**

`libc`

### **Description:**

The `iofunc_open_default()` function implements the default actions for the `_IO_CONNECT` message in a resource manager. This function calls:

- `iofunc_open()` to check the client's open mode against the resources attributes to see if the client can open the resource in that mode
- `iofunc_ocb_alloc()` to allocate an Open Control Block (OCB)
- `iofunc_ocb_attach()` to initialize the OCB
- `resmgr_open_bind()` to bind the newly-created OCB to the request.

You can place this function directly into the *connect\_funcs* table passed to *resmgr\_attach()*, at the *open* position, or you can call *iofunc\_func\_init()* to initialize all of the functions to their default values.

See the “Examples” section in the description of *iofunc\_open()* for the skeleton outline of the functionality (the second example, where *resmgr\_attach()* doesn’t specify that the managed resource is a directory).

## Returns:

|        |                                                                                                            |
|--------|------------------------------------------------------------------------------------------------------------|
| EOK    | Successful completion.                                                                                     |
| ENOSPC | There’s insufficient memory to allocate the OCB.                                                           |
| ENOMEM | There’s insufficient memory to allocate an internal data structure required by <i>resmgr_open_bind()</i> . |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_ocb\_attach()*, *iofunc\_ocb\_calloc()*,  
*iofunc\_open()*, *iofunc\_time\_update()*, *resmgr\_attach()*,  
*resmgr\_connect\_funcs\_t*, *resmgr\_open\_bind()*

Writing a Resource Manager chapter of the *Programmer’s Guide*.



## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_openfd(resmgr_context_t *ctp,
 io_openfd_t *msg,
 iofunc_ocb_t *ocb,
 iofunc_attr_t *attr);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_openfd_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

## Library:

`libc`

## Description:

The `iofunc_openfd()` helper function examines the mode specified by the `_IO_OPENFD` message, and increments the read and write count flags (`ocb->attr->rcount` and `ocb->attr->>wcount`), and the locking flags (`ocb->attr->rlocks` and `ocb->attr->>wlocks`), as specified by the open mode.

The function does what's needed to support the `openfd()` function.

## io\_openfd\_t structure

The `io_openfd_t` structure holds the `_IO_OPENFD` message received by the resource manager:

```
struct _io_openfd {
 uint16_t type;
 uint16_t combine_len;
 uint32_t ioflag;
 uint16_t sflag;
 uint16_t xtype;
 struct _msg_info info;
 uint32_t reserved2;
 uint32_t key;
};

typedef union {
 struct _io_openfd i;
} io_openfd_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_openfd` that contains the following members:

|                    |                                                                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_OPENFD</code> .                                                                                                                                                                                                                                                                                                                 |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> .                                                                                                                       |
| <i>ioflag</i>      | How the client wants to open the file; a combination of the following bits: <ul style="list-style-type: none"><li>● <code>O_RDONLY</code> — permit the file to be only read.</li><li>● <code>O_WRONLY</code> — permit the file to be only written.</li><li>● <code>O_RDWR</code> — permit the file to be both read and written.</li></ul> |

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              | <ul style="list-style-type: none"> <li>● O_APPEND — cause each record that's written to be written at the end of the file.</li> <li>● O_TRUNC — if the file exists, truncate it to contain no data. This flag has no effect if the file doesn't exist.</li> </ul>                                                                                                                                                                       |
| <i>sflag</i> | <p>How the client wants the file to be shared; a combination of the following bits:</p> <ul style="list-style-type: none"> <li>● SH_COMPAT — set compatibility mode.</li> <li>● SH_DENYRW — prevent read or write access to the file.</li> <li>● SH_DENYWR — prevent write access to the file.</li> <li>● SH_DENYRD — prevent read access to the file.</li> <li>● SH_DENYNO — permit both read and write access to the file.</li> </ul> |
| <i>xtype</i> | <p>Extended type information that can change the behavior of an I/O function. One of:</p> <ul style="list-style-type: none"> <li>● _IO_OPENFD_NONE — no extended type information.</li> <li>● _IO_OPENFD_PIPE — a pipe is being opened.</li> <li>● _IO_OPENFD_RESERVED — reserved</li> </ul>                                                                                                                                            |
| <i>info</i>  | <p>A pointer to a <code>_msg_info</code> structure that contains information about the message received by the resource manager.</p>                                                                                                                                                                                                                                                                                                    |
| <i>key</i>   | <p>Reserved for future use.</p>                                                                                                                                                                                                                                                                                                                                                                                                         |

**Returns:**

|        |                                             |
|--------|---------------------------------------------|
| EOK    | Success.                                    |
| EACCES | You don't have permission to open the file. |
| EBUSY  | The file has shared locks that are in use.  |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_ocb\_attach()*, *iofunc\_openfd\_default()*, `_msg_info`, *openfd()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_openfd_default(resmgr_context_t *ctp,
 io_openfd_t *msg,
 iofunc_ocb_t *ocb);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_openfd_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_openfd()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

## Library:

`libc`

## Description:

The function `iofunc_openfd_default()` function implements POSIX semantics for the client's `openfd()` call, which is received as an `_IO_OPENFD` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `openfd` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_openfd_default()` function calls `iofunc_openfd()` to do the actual work, and (if installed in the `io_funcs` table) issues the reply back to the client.

**Returns:**

- EOK            Success.
- EACCES        You don't have permission to open the file.
- EBUSY        The file has shared locks that are in use.
- EINVAL        The message type is invalid.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                Yes

**See also:**

*iofunc\_chown\_default()*, *iofunc\_func\_init()*, **iofunc\_ocb\_t**,  
*iofunc\_openfd()*, *iofunc\_sync\_default()*, *resmgr\_attach()*,  
**resmgr\_context\_t**, **resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_pathconf(resmgr_context_t *ctp,
 io_pathconf_t *msg,
 iofunc_ocr_t *ocr,
 iofunc_attr_t *attr);
```

**Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_pathconf_t` structure that contains the message that the resource manager received; see below.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

**Library:**

`libc`

**Description:**

The `iofunc_pathconf()` helper function does what's needed to support `pathconf()` with the `mount` and `attr` passed to it. Other `fsys_pathconf()` requests need to be handled by the caller.

If you write your own `pathconf` callout for your resource manager, use the following macro to pass the requested value back to the caller:

```
_IO_SET_PATHCONF_VALUE(resmgr_context_t *ctp,
 int value)
```

## io\_pathconf\_t structure

The `io_pathconf_t` structure holds the `_IO_PATHCONF` message received by the resource manager:

```
struct _io_pathconf {
 uint16_t type;
 uint16x_t combine_len;
 short name;
 uint16_t zero;
};

typedef union {
 struct _io_pathconf i;
 /* value is returned with MsgReply */
} io_pathconf_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_pathconf` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_PATHCONF</code> .                                                                                                                                                                                         |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |
| <i>name</i>        | The name of the configurable limit; see <i>pathconf()</i> .                                                                                                                                                         |

## Returns:

EOK, or `_RESMGR_DEFAULT` if the function didn't handle the *pathconf()* request.



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

`iofunc_attr_t`, `iofunc_ocb_t`, `iofunc_pathconf_default()`,  
`pathconf()`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_pathconf\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_PATHCONF messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_pathconf_default(resmgr_context_t *ctp,
 io_pathconf_t *msg,
 iofunc_ocr_t *ocr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_pathconf_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_pathconf()`.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_pathconf_default()` function implements POSIX semantics for the client's `pathconf()` call, which is received as an `_IO_PATHCONF` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `pathconf` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_pathconf_default()` function returns information about the resource, as per the POSIX specifications for `pathconf()`. The `iofunc_pathconf_default()` function simply calls `iofunc_pathconf()`, which does the actual work.

## Returns:

|        |                                                                                                                                                      |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| EOK    | Successful completion.                                                                                                                               |
| EINVAL | The pathconf parameter being ascertained wasn't one of <code>_PC_CHOWN_RESTRICTED</code> , <code>_PC_NO_TRUNC</code> , or <code>_PC_SYNC_IO</code> . |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_ocb\_t*, *iofunc\_pathconf()*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_read\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for IO\_READ messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_read_default(resmgr_context_t *ctp,
 io_read_t *msg,
 iofunc_ocr_t *ocr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_read_t` structure that contains the message that the resource manager received. For more information, see the documentation for `iofunc_read_verify()`.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_read_default()` function implements POSIX semantics for the client's `read()` call, which is received as an `IO_READ` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `read` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_read_default()` function calls `iofunc_read_verify()` to do the actual work, and (if installed in the `io_funcs` table) issues the reply back to the client.

## Returns:

- EBADF      The client doesn't have read access to this resource.
- EINVAL     The extended type information is invalid.
- EOK        The client has read access to this resource.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_ocb\_t*, *iofunc\_read\_verify()*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_read\_verify()***

© 2004, QNX Software Systems Ltd.

*Verify a client's read access to a resource*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_read_verify(resmgr_context_t* ctp,
 io_read_t* msg,
 iofunc_ocb_t* ocb,
 int* nonblock);
```

### **Arguments:**

|                 |                                                                                                                                                                                                                                                                                                              |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <b>resmgr_context_t</b> structure that the resource-manager library uses to pass context information between functions.                                                                                                                                                                       |
| <i>msg</i>      | A pointer to the <b>io_read_t</b> structure that contains the message that the resource manager received; see below.                                                                                                                                                                                         |
| <i>ocb</i>      | A pointer to the <b>iofunc_ocb_t</b> structure for the Open Control Block that was created when the client opened the resource.                                                                                                                                                                              |
| <i>nonblock</i> | NULL, or a pointer to a location where the function can store a value that indicates whether or not the device is nonblocking: <ul style="list-style-type: none"><li>• Zero — the client doesn't want to be blocked (i.e. O_NONBLOCK was set).</li><li>• Nonzero — the client wants to be blocked.</li></ul> |

### **Library:**

**libc**

### **Description:**

The *iofunc\_read\_verify()* helper function checks that the client that sent the **\_IO\_READ** message actually has read access to the resource, and, if *nonblock* isn't NULL, sets *nonblock* to **O\_NONBLOCK** or **0**).

The read permission check is done against *ocb->ioflag*.

Note that the `io_read_t` message has an override flag called *msg->i.xtype*. This flag allows the client to override the default blocking behavior for the resource on a per-request basis. This override flag is checked, and returned in the optional *nonblock*.

Note that if you're reading from a directory entry, you must return `struct dirent` structures in the *read* callout for your resource manager.

You'll also need to indicate how many bytes were read. You can do this with the macro:

```
_IO_SET_READ_NBYTES(resmgr_context_t *ctp,
 int nbytes)
```

### `io_read_t` structure

The `io_read_t` structure holds the `_IO_READ` message received by the resource manager:

```
struct _io_read {
 uint16_t type;
 uint16_t combine_len;
 int32_t nbytes;
 uint32_t xtype;
 uint32_t zero;
};

typedef union {
 struct _io_read i;
 /* unsigned char data[nbytes]; */
 /* nbytes is returned with MsgReply */
} io_read_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_read` that contains the following members:

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | _IO.READ.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <i>combine_len</i> | If the message is a combine message, _IO.COMBINE.FLAG is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <i>nbytes</i>      | The number of bytes that the client wants to read.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>xtype</i>       | Extended type information; one of: <ul style="list-style-type: none"><li>• _IO.XTYPE.NONE</li><li>• _IO.XTYPE.READCOND</li><li>• _IO.XTYPE.MQUEUE</li><li>• _IO.XTYPE.TCPIP</li><li>• _IO.XTYPE.TCPIP_MSG</li><li>• _IO.XTYPE.OFFSET</li><li>• _IO.XTYPE.REGISTRY</li><li>• _IO.XFLAG.DIR.EXTRA.HINT — this flag is valid only when reading from a directory. The filesystem should normally return extra directory information when it’s easy to get. If this flag is set, it is a hint to the filesystem to try harder (possibly causing media lookups) to return the extra information. The most common use would be to return _DTYPE.LSTAT information.</li><li>• _IO.XFLAG.NONBLOCK</li><li>• _IO.XFLAG.BLOCK</li></ul> |

For more information, see “Handling other read/write details” in the Writing a Resource Manager chapter of the *Programmer’s Guide*.

The commented-out declaration for *data* indicates that *nbytes* bytes of data immediately follow the `io_read_t` structure.



**Returns:**

- EOK        The client has read access to this resource.
- EBADF     The client doesn't have read access to this resource.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_open()*, *iofunc\_write\_verify()*

Writing a Resource Manager chapter of the *Programmer's Guide*

## ***iofunc\_readlink()***

© 2004, QNX Software Systems Ltd.

*Verify a client's ability to read a symbolic link*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_readlink(resmgr_context_t *ctp,
 io_readlink_t *msg,
 iofunc_attr_t *attr,
 struct _client_info *info);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_readlink_t` structure that contains the message that the resource manager received; see below.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.
- info* A pointer to a `_client_info` structure that contains the information about a client connection. For information about this structure, see *ConnectClientInfo()*.

### **Library:**

`libc`

### **Description:**

The *iofunc\_readlink()* helper function supports *readlink()* requests by verifying that the client can read a symbolic link. It's similar to *iofunc\_open()*.

The *iofunc\_read()* function checks to see if the client (described by the optional *info* structure) has access to open the resource (name passed in the *msg* structure). The *attr* structure describes the resource's attributes.

The *info* argument can be passed as NULL, in which case *iofunc\_read()* obtains the client information itself via a call to *iofunc\_client\_info()*. It is, of course, more efficient to get the client info once, rather than calling this function with NULL every time.

The *iofunc\_readlink()* function handles the readlink verification for the POSIX layer.

### **io\_readlink\_t structure**

The *io\_readlink\_t* structure holds the `_IO_CONNECT` message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_fstype_reply fstype_reply;
} io_readlink_t;
```

This message structure is a union of an input message (coming to the resource manager), `_io_connect`, and two possible output or reply messages (going back to the client):

- `_io_connect_link_reply` if the reply is redirecting the client to another resource
- Or:
- `_io_connect_fstype_reply` if the reply consists of a status and a file type.

### **Returns:**

|          |                                  |
|----------|----------------------------------|
| EBADFSYS | NULL was passed in <i>attr</i> . |
| EOK      | Successful completion.           |

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`_io_connect`, `_io_connect_link_reply`,  
`_io_connect_ftype_reply`, `iofunc_open()`, `readlink()`

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_rename(resmgr_context_t* ctp,
 io_rename_t* msg,
 iofunc_attr_t* oldattr,
 iofunc_attr_t* olddattr,
 iofunc_attr_t* newattr,
 iofunc_attr_t* newdattr,
 struct _client_info* info);
```

**Arguments:**

|                 |                                                                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <b>resmgr_context_t</b> structure that the resource-manager library uses to pass context information between functions.                                                |
| <i>msg</i>      | A pointer to the <b>io_rename_t</b> structure that contains the message that the resource manager received; see below.                                                                |
| <i>oldattr</i>  | A pointer to the <b>iofunc_attr_t</b> structure that describes the characteristics of the resource.                                                                                   |
| <i>olddattr</i> | NULL, or a pointer to the <b>iofunc_attr_t</b> structure that describes the characteristics of the parent directory.                                                                  |
| <i>newattr</i>  | A pointer to the <b>iofunc_attr_t</b> structure that describes the characteristics of the target, if it exists.                                                                       |
| <i>newdattr</i> | NULL, or a pointer to the <b>iofunc_attr_t</b> structure that describes the characteristics of the parent directory of the target.                                                    |
| <i>info</i>     | NULL, or a pointer to a <b>_client_info</b> structure that contains the information about a client connection. For information about this structure, see <i>ConnectClientInfo()</i> . |

## Library:

libc

## Description:

The function *iofunc\_rename()* does permission checks for the `_IO_CONNECT` message (subtype `_IO_CONNECT_RENAME`) for context *ctp*. The *newattr* argument is the attribute of the target if it already exists.

This function is similar to *iofunc\_open()*. The *iofunc\_rename()* function checks to see if the client (described by the optional *info* structure) has access to open the resource (name passed in the *msg* structure). The *attr* structure describes the resource's attributes.

The *info* argument can be passed as NULL, in which case *iofunc\_rename()* obtains the client information itself via a call to *iofunc\_client\_info()*. It is, of course, more efficient to get the client information once, rather than call this function with NULL every time.

## `io_rename_t` structure

The `io_rename_t` structure holds the `_IO_CONNECT` message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_ftype_reply ftype_reply;
} io_rename_t;
```

This message structure is a union of an input message (coming to the resource manager), `_io_connect`, and two possible output or reply messages (going back to the client):

- `_io_connect_link_reply` if the reply is redirecting the client to another resource

Or:

- `_io_connect_ftype_reply` if the reply consists of a status and a file type.

The reply includes the following extra information:

```
typedef union _io_rename_extra {
 char path[1];
} io_rename_extra_t;
```

## Returns:

|           |                                                                                                                                                       |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES    | The client doesn't have permissions to do the operation.                                                                                              |
| EBADFSYS  | NULL was passed in <i>oldattr</i> , <i>olddattr</i> , or <i>newdattr</i> .                                                                            |
| EFAULT    | A fault occurred when the kernel tried to access the <i>info</i> buffer.                                                                              |
| EINVAL    | The <i>oldattr</i> and <i>newdattr</i> have identical values, the client process is no longer valid, or attempt to remove the parent (".") directory. |
| EISDIR    | The old link is a directory but the new link isn't a directory.                                                                                       |
| ENOTDIR   | Attempt to unlink a nondirectory entry using directory semantics (e.g. <code>rmdir file</code> ).                                                     |
| ENOTEMPTY | Attempt to remove a directory that isn't empty.                                                                                                       |
| EOK       | Successful completion or there was already a <i>newattr</i> entry.                                                                                    |
| EPERM     | The group ID or owner ID didn't match.                                                                                                                |
| EROFS     | Attempt to remove an entry on a read-only filesystem.                                                                                                 |

## Classification:

QNX Neutrino

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

`_io_connect`, `_io_connect_link_reply`,  
`_io_connect_fstype_reply`, `iofunc_client_info()`, `iofunc_open()`

Writing a Resource Manager chapter of the *Programmer's Guide*.



### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_space_verify(resmgr_context_t *ctp,
 io_space_t *msg,
 iofunc_ocb_t *ocb,
 int *nonblock);
```

### **Arguments:**

- |                 |                                                                                                                                                                                                                                                                                                              |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions.                                                                                                                                                                 |
| <i>msg</i>      | A pointer to the <code>io_space_t</code> structure that contains the message that the resource manager received; see below.                                                                                                                                                                                  |
| <i>ocb</i>      | A pointer to the <code>iofunc_ocb_t</code> structure for the Open Control Block that was created when the client opened the resource.                                                                                                                                                                        |
| <i>nonblock</i> | NULL, or a pointer to a location where the function can store a value that indicates whether or not the device is nonblocking: <ul style="list-style-type: none"><li>• Zero — the client doesn't want to be blocked (i.e. O_NONBLOCK was set).</li><li>• Nonzero — the client wants to be blocked.</li></ul> |

### **Library:**

`libc`

### **Description:**

The `iofunc_space_verify()` helper function checks the client's permission for an `_IO_SPACE` message.

**io\_space\_t structure**

The `io_space_t` structure holds the `_IO_SPACE` message received by the resource manager:

```

struct _io_space {
 uint16_t type;
 uint16_t combine_len;
 uint16_t subtype;
 short whence;
 uint64_t start;
 uint64_t len;
};

typedef union {
 struct _io_space i;
 uint64_t o;
} io_space_t;

```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The `i` member is a structure of type `_io_space` that contains the following members:

|                    |                                                                                                                                                                                                                                                                                                              |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_SPACE</code> .                                                                                                                                                                                                                                                                                     |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> .                                                                                          |
| <i>subtype</i>     | <code>F_ALLOCSP</code> or <code>F_FREESP</code> .                                                                                                                                                                                                                                                            |
| <i>whence</i>      | The position in the file. The possible values (defined in <code>&lt;unistd.h&gt;</code> ) are: <ul style="list-style-type: none"> <li><code>SEEK_CUR</code> The new file position is computed relative to the current file position. The value of <i>start</i> may be positive, negative or zero.</li> </ul> |

SEEK\_END The new file position is computed relative to the end of the file.

SEEK\_SET The new file position is computed relative to the start of the file. The value of *start* must not be negative.

*start* The relative offset from the file position determined by the *whence* member.

*len* The relative size by which to increase the file.  
A value of zero means to end of file.

The *o* member is the file size.

### Returns:

EBADF The client doesn't have read access to this resource.

EISDIR The resource is a directory.

EOK The client has read access to this resource.

### Classification:

QNX Neutrino

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_ocb_t`, `iofunc_open()`, `iofunc_write_default()`,  
`iofunc_write_verify()`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_stat(resmgr_context_t* ctp,
 iofunc_attr_t* attr,
 struct stat* stat);
```

## Arguments:

- ctp* A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions.
- attr* A pointer to the **iofunc\_attr\_t** structure that describes the characteristics of the device that's associated with your resource manager.
- stat* A pointer to the **stat** structure that you want to fill. For more information, see *stat()*.

## Library:

**libc**

## Description:

The *iofunc\_stat()* function populates the passed *stat* structure based on information from the passed *attr* structure and the context pointer, *ctp*.

This is typically used when the resource manager is handling the **\_IO\_STAT** message, and needs to format the current status information for the resource.

## Returns:

**EOK** Successful completion.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_attr_t`, `iofunc_stat_default()`, `iofunc_time_update()`,  
`resmgr_context_t`, `stat()`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_stat_default(resmgr_context_t *ctp,
 io_stat_t *msg,
 iofunc_ocb_t *ocb);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_stat_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

## Library:

`libc`

## Description:

The `iofunc_stat_default()` function implements POSIX semantics for the client's `stat()` or `fstat()` call, which is received as an `_JO_STAT` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `stat` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_stat_default()` function calls:

- `iofunc_time_update()`, to ensure that the time entries in the `ocb->attr` structure are current and valid
- `iofunc_stat()` to construct a status entry based on the information in the `ocb->attr` structure.

## **io\_stat\_t structure**

The `io_stat_t` structure holds the `_IO_STAT` message received by the resource manager:

```
struct _io_stat {
 uint16_t type;
 uint16_t combine_len;
 uint32_t zero;
};

typedef union {
 struct _io_stat i;
 struct stat o;
} io_stat_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client).

The `i` member is a structure of type `_io_stat` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_STAT</code> .                                                                                                                                                                                             |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer’s Guide</i> . |

The `o` member is a structure of type `stat`; for more information, see `stat()`.

## **Returns:**

EOK      Successful completion.



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_ocb\_t*, *iofunc\_stat()*,  
*iofunc\_time\_update()*, *resmgr\_attach()*, *resmgr\_context\_t*,  
*resmgr\_io\_funcs\_t*, *stat()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_sync()***

© 2004, QNX Software Systems Ltd.

*Indicate if synchronization is needed*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_sync(resmgr_context_t* ctp,
 iofunc_ocb_t* ocb,
 int ioflag);
```

### **Arguments:**

- ctp*      A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- ocb*      A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- ioflag*    The operation being performed on the OCB:
- `_IO_FLAG_WR` — writing.
  - `_IO_FLAG_RD` — reading.

### **Library:**

`libc`

### **Description:**

The `iofunc_sync()` function indicates if some form of synchronization is needed.

### **Returns:**

- `O_DSYNC`    Data integrity is needed.
- `O_SYNC`     File integrity is needed.
- `0`           Synchronization isn't needed.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_open()*, *iofunc\_write\_default()*, *iofunc\_write\_verify()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_sync\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_SYNC messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_sync_default(resmgr_context_t *ctp,
 io_sync_t *msg,
 iofunc_ocr_t *ocr);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_sync_t` structure that contains the message that the resource manager received. For more information, see `iofunc_sync_verify()`.
- ocr* A pointer to the `iofunc_ocr_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The function `iofunc_sync_default()` function implements POSIX semantics for the client's `sync()` call, which is received as an `_IO_SYNC` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `sync` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_sync_default()` function calls `iofunc_sync_verify()` which checks to see if the client can synchronize the resource, and (if installed in the `io_funcs` table) issues the reply back to the client.

## Returns:

- EINVAL     The resource doesn't support synchronizing.
- EOK        The client can synchronize the resource.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*iofunc\_func\_init()*, *iofunc\_ocb\_t*, *iofunc\_sync()*,  
*iofunc\_sync\_verify()*, *resmgr\_attach()*, *resmgr\_context\_t*,  
*resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_sync\_verify()***

© 2004, QNX Software Systems Ltd.

*Verify permissions to sync*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_sync_verify(resmgr_context_t *ctp,
 io_sync_t *msg,
 iofunc_ocb_t *ocb);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*msg* A pointer to the `io_sync_t` structure that contains the message that the resource manager received; see below.

*ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_sync_verify()` function verifies that the client has permission to synchronize.

### **io\_sync\_t structure**

The `io_sync_t` structure holds the `_IO_SYNC` message received by the resource manager:

```
struct _io_sync {
 uint16_t type;
 uint16_t combine_len;
 uint32_t flag;
};
```

```
typedef union {
 struct _io_sync i;
} io_sync_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_sync` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_SYNC</code> .                                                                                                                                                                                             |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |
| <i>flag</i>        | One of: <ul style="list-style-type: none"><li>• <code>O_DSYNC</code></li><li>• <code>O_RSynchronize</code></li><li>• <code>O_SYNC</code></li></ul>                                                                  |

For more information about these flags, see *open()*.

## Returns:

|                     |                                              |
|---------------------|----------------------------------------------|
| <code>EINVAL</code> | The resource doesn't support syncing.        |
| <code>EOK</code>    | The client has read access to this resource. |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_ocb_t`, `iofunc_open()`, `iofunc_write_default()`,  
`iofunc_write_verify()`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.



### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_time_update(iofunc_attr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

### **Library:**

`libc`

### **Description:**

The `iofunc_time_update()` function examines the *flags* member in the passed *attr* structure against the bits `IOFUNC_ATTR_ETIME`, `IOFUNC_ATTR_MTIME`, and `IOFUNC_ATTR_CTIME`. If any of these bits are set, the corresponding time member of *attr* (e.g. *attr->etime*) isn't valid. This function updates all invalid *attr* members to the current time.

If `iofunc_time_update()` makes any change to the *attr* structure's time members, it sets `IOFUNC_ATTR_DIRTY_TIME` in the *attr* structure's *flags* member. This function always clears the `IOFUNC_ATTR_ETIME`, `IOFUNC_ATTR_MTIME`, and `IOFUNC_ATTR_CTIME` bits from *attr->flags*.

### **Returns:**

`EOK` Successful completion.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`iofunc_attr_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_unblock(resmgr_context_t * ctp,
 iofunc_attr_t * attr);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

## Library:

`libc`

## Description:

The `iofunc_unblock()` function unblocks any clients that are blocked on any internal resource manager structures.



Currently, this involves only the advisory lock list that's maintained by the attribute.

---

If a client connection is found:

- that client is unblocked, and is replied to with the error `EINTR`.
- `iofunc_unblock()` returns `_RESMGR_NOREPLY`.

If no client connection is found, `iofunc_unblock()` returns `_RESMGR_DEFAULT`.

**Returns:**

`_RESMGR_DEFAULT`

No client connection was found.

`_RESMGR_NOREPLY`

A client connection has been unblocked.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    No

Interrupt handler    No

Signal handler    Yes

Thread    Yes

**See also:**

*iofunc\_unblock\_default()*

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_unblock_default(resmgr_context_t * ctp,
 io_pulse_t * msg,
 iofunc_ocb_t * ocb);
```

**Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_pulse_t` structure that describes the pulse that the resource manager received.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

**Library:**

`libc`

**Description:**

The `iofunc_unblock_default()` function calls `iofunc_unblock()`.

The `iofunc_unblock_default()` function implements the functionality required when the client requests to be unblocked (e.g. a signal or timeout).

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `unblock` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The unblock message is synthesized by the resource-manager shared library when a client wishes to unblock from its `MsgSendv()` to the

resource manager. The *iofunc\_unblock\_default()* function takes care of freeing up any locks that the client may have placed on the resource.

## Returns:

`_RESMGR_DEFAULT`

No client connection was found.

`_RESMGR_NOREPLY`

A client connection has been unblocked.

## Examples:

If you're calling *iofunc\_lock\_default()*, your unblock handler should call *iofunc\_unblock\_default()*:

```
if((status = iofunc_unblock_default(...)) != _RESMGR_DEFAULT) {
 return status;
}

/* Do your own thing to look for a client to unblock */
```

## Classification:

QNX Neutrino

### Safety

---

Cancellation point    No

Interrupt handler    No

Signal handler    Yes

Thread    Yes

**See also:**

*iofunc\_func\_init()*, *iofunc\_lock\_default()*, *iofunc\_ocb\_t*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_unlink()***

© 2004, QNX Software Systems Ltd.

*Verify that an entry can be unlinked*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_unlink(resmgr_context_t* ctp,
 io_unlink_t* msg,
 iofunc_attr_t* attr,
 iofunc_attr_t* dattr,
 struct _client_info* info);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_unlink_t` structure that contains the message that the resource manager received; see below.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the resource.
- dattr* NULL, or a pointer to the `iofunc_attr_t` structure that describes the characteristics of the parent directory.
- info* NULL, or a pointer to a `_client_info` structure that contains information about the client. For information about this structure, see `ConnectClientInfo()`.

### **Library:**

`libc`

### **Description:**

The `iofunc_unlink()` function verifies that the `msg` specifies valid semantics for an unlink, and that the client is allowed to unlink the resource, as specified by a combination of who the client is (`info`), and the resource attributes `attr`, `dattr`, `attr->uid` and `attr->gid`.



If a directory entry is being removed, *iofunc\_unlink()* checks to see that the directory is empty. The *iofunc\_unlink()* function also updates the time stamps, and decrements the link count for the entry.

### **io\_unlink\_t structure**

The `io_unlink_t` structure holds the `_IO_CONNECT` message received by the resource manager:

```
typedef union {
 struct _io_connect connect;
 struct _io_connect_link_reply link_reply;
 struct _io_connect_fstype_reply fstype_reply;
} io_unlink_t;
```

This message structure is a union of an input message (coming to the resource manager), `_io_connect`, and two possible output or reply messages (going back to the client):

- `_io_connect_link_reply` if the reply is redirecting the client to another resource
- Or:
- `_io_connect_fstype_reply` if the reply consists of a status and a file type.

### **Returns:**

|           |                                                                                                    |
|-----------|----------------------------------------------------------------------------------------------------|
| EOK       | Successful completion.                                                                             |
| ENOTDIR   | Attempt to unlink a nondirectory entry using directory semantics, (e.g. <code>rmdir file</code> ). |
| EINVAL    | Attempt to remove the <code>."</code> directory.                                                   |
| ENOTEMPTY | Attempt to remove a directory that isn't empty.                                                    |
| EROFS     | Attempt to remove an entry on a read-only filesystem.                                              |
| EACCES    | The client doesn't have permissions to do the operation.                                           |
| EPERM     | The group ID or owner ID didn't match.                                                             |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`_io_connect`, `_io_connect_link_reply`,  
`_io_connect_fstype_reply`, `ConnectClientInfo()`,  
`iofunc_attr_t`, `iofunc_check_access()`, `resmgr_context_t`

Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_unlock\_ocb\_default()***

*Default handler for the unlock\_ocb callout*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_unlock_ocb_default(
 resmgr_context_t * ctp,
 void * reserved,
 iofunc_ocb_t * ocb);
```

### **Arguments:**

|                 |                                                                                                                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <b>resmgr_context_t</b> structure that the resource-manager library uses to pass context information between functions. |
| <i>reserved</i> | This argument must be NULL.                                                                                                            |
| <i>ocb</i>      | A pointer to the <b>iofunc_ocb_t</b> structure for the Open Control Block that was created when the client opened the resource.        |

### **Library:**

**libc**

### **Description:**

The *iofunc\_unlock\_ocb\_default()* function calls *iofunc\_attr\_unlock()* to enforce unlocking on the attributes for the group of messages that were sent by the client.

You can place this function directly into the *io\_funcs* table passed to *resmgr\_attach()*, at the *unlock\_ocb* position, or you can call *iofunc\_func\_init()* to initialize all of the functions to their default values.

**Returns:**

- EOK        Success.
- EAGAIN     On the first use, all kernel mutex objects were in use.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*iofunc\_attr\_unlock()*, *iofunc\_func\_init()*, *iofunc\_ocb\_t*,  
*resmgr\_attach()*, *resmgr\_context\_t*, *resmgr\_io\_funcs\_t*

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_utime(resmgr_context_t* ctp,
 io_utime_t* msg,
 iofunc_ocb_t* ocb,
 iofunc_attr_t* attr);
```

**Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_utime_t` structure that contains the message that the resource manager received; see below.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.
- attr* A pointer to the `iofunc_attr_t` structure that describes the characteristics of the device that's associated with your resource manager.

**Library:**

`libc`

**Description:**

The `iofunc_utime()` helper function examines the `flags` member in the passed `attr` structure and sets the `IOFUNC_ATTR_ETIME` and `IOFUNC_ATTR_MTIME` bits if requested.

The function sets the `IOFUNC_ATTR_CTIME` and `IOFUNC_ATTR_DIRTY_TIME` bits. It then calls `iofunc_time_update()` to update the file times.

## io\_utime\_t structure

The `io_utime_t` structure holds the `_IO_ETIME` message received by the resource manager:

```
struct _io_utime {
 uint16_t type;
 uint16_t combine_len;
 int32_t cur_flag;
 struct utimbuf times;
};

typedef union {
 struct _io_utime i;
} io_utime_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_utime` that contains the following members:

|                    |                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>        | <code>_IO_ETIME</code> .                                                                                                                                                                                            |
| <i>combine_len</i> | If the message is a combine message, <code>_IO_COMBINE_FLAG</code> is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the <i>Programmer's Guide</i> . |
| <i>cur_flag</i>    | If set, <code>iofunc_utime()</code> ignores the <i>times</i> member, and set the appropriate file times to the current time.                                                                                        |
| <i>times</i>       | A <code>utimbuf</code> structure that specifies the time to use when setting the file times. For more information about this structure, see <code>utime()</code> .                                                  |

**Returns:**

|        |                                                                          |
|--------|--------------------------------------------------------------------------|
| EACCES | The client doesn't have permissions to do the operation.                 |
| EFAULT | A fault occurred when the kernel tried to access the <i>info</i> buffer. |
| EINVAL | The client process is no longer valid.                                   |
| ENOSYS | NULL was passed in <i>info</i> .                                         |
| EOK    | Successful completion.                                                   |
| EPERM  | The group ID or owner ID didn't match.                                   |
| EROFS  | Attempt to remove an entry on a read-only filesystem.                    |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***iofunc\_time\_update()*, *iofunc\_utime\_default()*, *utime()*Writing a Resource Manager chapter of the *Programmer's Guide*.

## ***iofunc\_utime\_default()***

© 2004, QNX Software Systems Ltd.

*Default handler for \_IO\_UTIME messages*

### **Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_utime_default(resmgr_context_t* ctp,
 io_utime_t* msg,
 iofunc_ocb_t* ocb);
```

### **Arguments:**

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_utime_t` structure that contains the message that the resource manager received; see `iofunc_utime()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

### **Library:**

`libc`

### **Description:**

The `iofunc_utime_default()` function implements POSIX semantics for the client's `utime()` call, which is received as an `_IO_UTIME` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `utime` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_utime_default()` function calls `iofunc_utime()`, which does the actual work. It verifies that the client has the necessary permissions to effect a `utime` on the device. If so, the `utime` is



performed, modifying elements of the *ocb->attr* structure. This function takes care of updating these bits in the flags member of *ocb->attr*:

- IOFUNC\_ATTR\_ETIME
- IOFUNC\_ATTR\_CTIME
- IOFUNC\_ATTR\_MTIME
- IOFUNC\_ATTR\_DIRTY\_TIME
- IOFUNC\_ATTR\_DIRTY\_MODE

The *iofunc\_utime()* function then calls *iofunc\_time\_update()* to update the appropriate time fields in *ocb->attr*.

## Returns:

|        |                                                          |
|--------|----------------------------------------------------------|
| EOK    | Successful completion.                                   |
| EROFS  | An attempt was made to utime a read-only filesystem.     |
| EACCES | The client doesn't have permissions to do the operation. |
| EPERM  | The group ID or owner ID didn't match.                   |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_func\_init()*, *iofunc\_time\_update()*, **iofunc\_ocb\_t**,  
*iofunc\_utime()*, *resmgr\_attach()*, **resmgr\_context\_t**,  
**resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

## Synopsis:

```
#include <sys/iofunc.h>

int iofunc_write_default(resmgr_context_t* ctp,
 io_write_t* msg,
 iofunc_ocb_t* ocb);
```

## Arguments:

- ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.
- msg* A pointer to the `io_write_t` structure that contains the message that the resource manager received. For more information, see `iofunc_write_verify()`.
- ocb* A pointer to the `iofunc_ocb_t` structure for the Open Control Block that was created when the client opened the resource.

## Library:

`libc`

## Description:

The `iofunc_write_default()` function implements POSIX semantics for the client's `write()` call, which is received as an `_IO_WRITE` message by the resource manager.

You can place this function directly into the `io_funcs` table passed to `resmgr_attach()`, at the `write` position, or you can call `iofunc_func_init()` to initialize all of the functions to their default values.

The `iofunc_write_default()` function calls `iofunc_write_verify()` to do the actual work, and (if installed in the `io_funcs` table) issues the reply back to the client.

**Returns:**

- EBADF      The client doesn't have read access to this resource.
- EINVAL     An unknown *xtype* was given.
- EOK        The client has read access to this resource.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                Yes

**See also:**

*iofunc\_func\_init()*, **iofunc\_ocb\_t**, *iofunc\_open()*,  
*iofunc\_write\_verify()*, *resmgr\_attach()*, **resmgr\_context\_t**,  
**resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the *Programmer's Guide*.

**Synopsis:**

```
#include <sys/iofunc.h>

int iofunc_write_verify(resmgr_context_t* ctp,
 io_write_t* msg,
 iofunc_ocr_t* ocb,
 int* nonblock);
```

**Arguments:**

- |                 |                                                                                                                                                                                                                                                                                                                           |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>      | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions.                                                                                                                                                                              |
| <i>msg</i>      | A pointer to the <code>io_write_t</code> structure that contains the message that the resource manager received; see below.                                                                                                                                                                                               |
| <i>ocr</i>      | A pointer to the <code>iofunc_ocr_t</code> structure for the Open Control Block that was created when the client opened the resource.                                                                                                                                                                                     |
| <i>nonblock</i> | NULL, or a pointer to a location where the function can store a value that indicates whether or not the device is nonblocking: <ul style="list-style-type: none"><li>• Zero — the client doesn't want to be blocked (i.e. <code>O_NONBLOCK</code> was set).</li><li>• Nonzero — the client wants to be blocked.</li></ul> |

**Library:**

`libc`

**Description:**

The `iofunc_write_verify()` function checks that the client that sent the write message actually has write access to the resource, and, optionally (if `nonblock` isn't NULL), sets `nonblock` to `O_NONBLOCK` or 0.

The write permission check is done against *ocb->ioflag*.

Note that the `io_write_t` message has an override flag called *msg->i.xtype*. This flag allows the client to override the default blocking behavior for the resource on a per-request basis. This override flag is checked, and returned in the optional *nonblock*.

In *write* callout for your resource manager, you'll need to indicate how many bytes were written. You can do this with the macro:

```
_IO_SET_WRITE_NBYTES(resmgr_context_t *ctp,
 int nbytes)
```

### io\_write\_t structure

The `io_write_t` structure holds the `_IO_WRITE` message received by the resource manager:

```
struct _io_write {
 uint16_t type;
 uint16_t combine_len;
 int32_t nbytes;
 uint32_t xtype;
 uint32_t zero;
 /* unsigned char data[nbytes]; */
};

typedef union {
 struct _io_write i;
 /* nbytes is returned with MsgReply */
} io_write_t;
```

The I/O message structures are unions of an input message (coming to the resource manager) and an output or reply message (going back to the client). In this case, there's only an input message, *i*.

The *i* member is a structure of type `_io_write` that contains the following members:

```
type _IO_WRITE.
```

- combine\_len* If the message is a combine message, `_IO_COMBINE_FLAG` is set in this member. For more information, see “Combine messages” in the Writing a Resource Manager chapter of the *Programmer’s Guide*.
- nbytes* The number of bytes that the client wants to write.
- xtype* Extended type information; one of:
- `_IO_XTYPE_NONE`
  - `_IO_XTYPE_READCOND`
  - `_IO_XTYPE_MQUEUE`
  - `_IO_XTYPE_TCPIP`
  - `_IO_XTYPE_TCPIP_MSG`
  - `_IO_XTYPE_OFFSET`
  - `_IO_XTYPE_REGISTRY`
  - `_IO_XFLAG_DIR_EXTRA_HINT` — this flag is valid only when reading from a directory. The filesystem should normally return extra directory information when it’s easy to get. If this flag is set, it is a hint to the filesystem to try harder (possibly causing media lookups) to return the extra information. The most common use would be to return `_DTYPE_LSTAT` information.
  - `_IO_XFLAG_NONBLOCK`
  - `_IO_XFLAG_BLOCK`

For more information, see “Handling other read/write details” in the Writing a Resource Manager chapter of the *Programmer’s Guide*.

The commented-out declaration for *data* indicates that *nbytes* bytes of data immediately follow the `_io_write` structure.

**Returns:**

- EOK        The client has write access to this resource.
- EBADF     The client doesn't have write access to this resource.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*iofunc\_read\_verify()*

Writing a Resource Manager chapter of the *Programmer's Guide*.



## Synopsis:

```
#include <unistd.h>
#include <sys/iomsg.h>

int ionotify (int fd,
 int action,
 int flags,
 const struct sigevent* event);
```

## Arguments:

- fd*            The file descriptor associated with the resource manager that you want to notify.
- action*        The type of arming action to take; see “Actions,” below.
- flags*         The types of conditions that can be checked for notification; see “Flags,” below.
- event*         A pointer to a **sigevent** structure that defines the event that you want the resource manager to send as a notification, or NULL to disarm a notification.

## Library:

**libc**

## Description:

The *ionotify()* function arms the resource manager associated with *fd* to send the event notification *event*. The event is sent when a condition specified by a combination of *action* and *flags* occurs.

## Flags

The *flags* argument specifies the types of conditions that can be checked for notification. Each resource manager maintains a different context for each notification condition. Only those notification bits specified are affected. In the following example, the second call to

*ionotify()* doesn't affect the first, since it specifies a different notification:

```
ionotify(fd, _NOTIFY_ACTION_POLLARM,
 _NOTIFY_COND_INPUT, &event);
ionotify(fd, _NOTIFY_ACTION_POLLARM,
 _NOTIFY_COND_OUTPUT, &event);
```

The conditions specified by *flags* are:

#### `_NOTIFY_COND_OBAND`

Out-of-band data is available. The definition of out-of-band data depends on the resource manager.

#### `_NOTIFY_COND_OUTPUT`

There's room in the output buffer for more data. The amount of room available needed to satisfy this condition depends on the resource manager. Some resource managers may default to an empty output buffer, while others may choose some percentage of the buffer empty.

#### `_NOTIFY_COND_INPUT`

There's input data available. The amount of data available defaults to 1. For a character device such as a serial port, this would be a character. For a POSIX message queue, it would be a message. Each resource manager selects an appropriate object.

The method for changing the default number for `_NOTIFY_COND_OUTPUT` and `_NOTIFY_COND_INPUT` depends on the device. For example, character special devices can call *readcond()*.

For resource managers that support both an edited and raw mode, the mode should be set to raw to ensure proper operation of *ionotify()*.

The above flags are located in the top bits of *flags*. They are defined by `_NOTIFY_COND_MASK`.

In the case of an asynchronous notification using the passed *event*, such as a Neutrino pulse or queued realtime signal, the 32-bit value in

`event->sigev_value.sival_int` is returned to you unmodified, *unless* you've selected the `SI_NOTIFY` code, in which case the top bits (defined by `_NOTIFY_COND_MASK`) are set to the active notifications. In this case, you should limit the *sival\_int* to the mask defined by `_NOTIFY_DATA_MASK`.

For example, the Unix `select()` function specifies `SI_NOTIFY` and uses the allowable data bits of *sival\_int* as a serial number.



If you're using the `SI_NOTIFY` code, then you should clear the bits as specified by `_NOTIFY_COND_MASK` in the *sigev\_value* field — the resource manager only ever ORs in a value, it never clears the bits.

---

## Actions

The *action* argument specifies the type of arming action to take. When a condition is armed, the resource manager monitors it and, when met, delivers *event* using `MsgDeliverEvent()`. When an event is delivered, it's always disarmed except where noted below.

Note that for transition arming (as specified by an *action* of `_NOTIFY_ACTION_TRANARM`, only *one* notification of that type can be outstanding *per device*. When the transition arm fires, it's removed.

Each action is designed to support a specific notification type as follows:

### `_NOTIFY_ACTION_POLL`

This action does a poll of the notification conditions specified by *flags*. It never arms an event, and it cancels all other asynchronous event notifications set up by a previous call to `ionotify()`. This also allows it to be used as a simple “disarm” call.

Returns active conditions as requested by *flags*.

### `_NOTIFY_ACTION_POLLARM`

This action does a poll in the same way as `_NOTIFY_ACTION_POLL`. However, if none of the conditions

specified in *flags* are present then each condition specified in *flags* is armed. If any condition is met, none of the conditions are armed. The Unix *select()* function uses *ionotify()* with this action.

Returns active conditions as requested by *flags*.

#### `_NOTIFY_ACTION_TRANARM`

This action arms for transitions of the notification conditions specified by *flags*. A transition is defined as a data transition from empty to nonempty on input. Its use on output isn't defined. Note that if there is data available when this call is used, a data transition *won't* occur. To generate an event using this type of notification, you must arm the event and then drain the input using a nonblocking read. After this point, new input data causes the event to be delivered. The *mq\_notify()* function uses *ionotify()* with this action.

Since this arms for a transition, the return value is always zero.

You can use the `_NOTIFY_ACTION_POLLARM` or `_NOTIFY_ACTION_POLL` action to generate events that are level- as opposed to transition-oriented.

When an action is armed in a resource manager, it remains armed until:

- A thread sets a new action (this disarms any current action and possibly arms a new action),
- The event is delivered and the action wasn't a continuous one,
- The thread closes the device.

### Returns:

Active conditions as requested by *flags*. In the case of a transition action, a zero is returned. If an error occurs, -1 is returned (*errno* is set).

**Errors:**

|            |                                                                                                                                                |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>fd</i> doesn't exist, or <i>fd</i> is no longer connected to a channel.                                         |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                     |
| EINTR      | The call was interrupted by a signal.                                                                                                          |
| ENOMEM     | The resource manager couldn't allocate a notify entry to save the request.                                                                     |
| ENOSYS     | The requested action isn't supported by this resource manager.                                                                                 |
| ESRVRFAULT | A fault occurred in a server's address space while accessing the server's message buffers. This may have occurred on the receive or the reply. |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                               |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`sigevent`

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(AF_INET,
 SOCK_RAW,
 proto);
```

**Description:**

IP is the transport layer protocol used by the Internet protocol family. You may set options at the IP level when you're using higher-level protocols based on IP, such as TCP and UDP. You may also access IP through a "raw socket" (when you're developing new protocols or special-purpose applications).

There are several IP-level *setsockopt()* and *getsockopt()* options. You can use `IP_OPTIONS` to provide IP options to be transmitted in the IP header of each outgoing packet or to examine the header options on incoming packets. IP options may be used with any socket type in the Internet family. The format of IP options to be sent is that specified by the IP protocol specification (*RFC-791*), with one exception: the list of addresses for Source Route options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address is extracted from the option list and the size adjusted accordingly before use. To disable previously specified options, use a zero-length buffer:

```
setsockopt(s, IPPROTO_IP, IP_OPTIONS, NULL, 0);
```

You can use `IP_TOS` and `IP_TTL` to set the type-of-service and time-to-live fields in the IP header for `SOCK_STREAM` and `SOCK_DGRAM` sockets. For example:

```
int tos = IPTOS_LOWDELAY; /* see <netinet/ip.h> */
setsockopt(s, IPPROTO_IP, IP_TOS, &tos, sizeof(tos));

int ttl = 60; /* max = 255 */
setsockopt(s, IPPROTO_IP, IP_TTL, &ttl, sizeof(ttl));
```

If the `IP_RECVDSTADDR` option is enabled on a `SOCK_DGRAM` or `SOCK_RAW` socket, the `recvmsg()` call returns the destination IP address for a UDP datagram. The `msg_control` field in the `msg_hdr` structure points to a buffer that contains a `cmsghdr` structure followed by the IP address. The `cmsghdr` fields have the following values:

```
msg_len = sizeof(struct cmsghdr) + sizeof(struct in_addr)
msg_level = IPPROTO_IP
msg_type = IP_RECVDSTADDR
```

If the `IP_RECVIF` option is enabled on a `SOCK_DGRAM` or `SOCK_RAW` socket, the `recvmsg()` call returns a `struct sockaddr_dl` corresponding to the interface on which the packet was received. The `msg_control` field in the `msg_hdr` structure points to a buffer that contains a `cmsghdr` structure followed by the `struct sockaddr_dl`. The `cmsghdr` fields have the following values:

```
msg_len = sizeof(struct cmsghdr) + sizeof(struct sockaddr_dl)
msg_level = IPPROTO_IP
msg_type = IP_RECVIF
```

Raw IP sockets are connectionless, and are normally used with the `sendto()` and `recvfrom()` calls, although you can also use `connect()` to fix the destination for future packets (in which case you can use the `read()` or `recv()` and `write()` or `send()` system calls).

If the `proto` parameter to `socket()` is 0, the default protocol `IPPROTO_RAW` is used for outgoing packets, and only incoming packets destined for that protocol are received. If `proto` is nonzero, that protocol number will be used on outgoing packets and to filter incoming packets.

Outgoing packets automatically have an IP header prepended to them (based on the destination address and the protocol number the socket is created with), unless the `IP_HDRINCL` option has been set.

Incoming packets are received with IP header and options intact.

`IP_HDRINCL` indicates the complete IP header is included with the data and may be used only with the `SOCK_RAW` type.



```
#include <netinet/ip.h>

int hinc1 = 1; /* 1 = on, 0 = off */
setsockopt(s, IPPROTO_IP, IP_HDRINCL, &hinc1, sizeof(hinc1));
```

The program must set all the fields of the IP header, including the following:

```
ip->ip_v = IPVERSION;
ip->ip_hl = hlen >> 2;
ip->ip_id = 0; /* 0 means kernel set appropriate value */
ip->ip_off = offset;
```

If the header source address is set to `INADDR_ANY`, the kernel chooses an appropriate address.

## Multicasting

IP multicasting is supported only on `AF_INET` sockets of type `SOCK_DGRAM` and `SOCK_RAW`, and only on networks where the interface driver supports multicasting.

## Multicast Options

### IP\_MULTICAST\_TTL

Change the time-to-live (TTL) for outgoing multicast datagrams in order to control the scope of the multicasts:

```
u_char ttl; /* range: 0 to 255, default = 1 */
setsockopt(s, IPPROTO_IP, IP_MULTICAST_TTL, &ttl, sizeof(ttl));
```

Datagrams with a TTL of 1 aren't forwarded beyond the local network. Multicast datagrams with a TTL of 0 aren't transmitted on any network, but may be delivered locally if the sending host belongs to the destination group and if multicast loopback hasn't been disabled on the sending socket (see below). Multicast datagrams with TTL greater than 1 may be forwarded to other networks if a multicast router is attached to the local network.

### IP\_MULTICAST\_IF

For hosts with multiple interfaces, each multicast transmission is sent from the primary network interface. The `IP_MULTICAST_IF` option overrides the default for subsequent transmissions from a given socket:

```
struct in_addr addr;
setsockopt(s, IPPROTO_IP, IP_MULTICAST_IF, &addr, sizeof(addr));
```

where *addr* is the local IP address of the desired interface or `INADDR_ANY` to specify the default interface. You can get an interface's local IP address and multicast capability by sending the `SIOCGIFCONF` and `SIOCGIFFLAGS` requests to `ioctl()`. Normal applications shouldn't need to use this option.

### IP\_MULTICAST\_LOOP

If a multicast datagram is sent to a group to which the sending host itself belongs (on the outgoing interface), a copy of the datagram is, by default, looped back by the IP layer for local delivery. The `IP_MULTICAST_LOOP` option gives the sender explicit control over whether or not subsequent datagrams are looped back:

```
u_char loop; /* 0 = disable, 1 = enable (default) */
setsockopt(s, IPPROTO_IP, IP_MULTICAST_LOOP, &loop, sizeof(loop));
```

This option improves performance for applications that may have no more than one instance on a single host (such as a router demon), by eliminating the overhead of receiving their own transmissions. It shouldn't generally be used by applications for which there may be more than one instance on a single host (such as a conferencing program) or for which the sender doesn't belong to the destination group (such as a time querying program).

A multicast datagram sent with an initial TTL greater than 1 may be delivered to the sending host on a different interface from that on which it was sent, if the host belongs to the

destination group on that other interface. The loopback control option has no effect on such delivery.

#### IP\_ADD\_MEMBERSHIP

A host must become a member of a multicast group before it can receive datagrams sent to the group. To join a multicast group, use the `IP_ADD_MEMBERSHIP` option:

```
struct ip_mreq mreq;
setsockopt(s, IPPROTO_IP, IP_ADD_MEMBERSHIP, &mreq, sizeof(mreq));
```

where *mreq* is the following structure:

```
struct ip_mreq {
 struct in_addr imr_multiaddr; /* multicast group to join */
 struct in_addr imr_interface; /* interface to join on */
}
```

Set *imr\_interface* to `INADDR_ANY` to choose the default multicast interface, or to the IP address of a particular multicast-capable interface if the host is multihomed. Membership is associated with a single interface; programs running on multihomed hosts may need to join the same group on more than one interface. Up to `IP_MAX_MEMBERSHIPS` (currently 20) memberships may be added on a single socket.

#### IP\_DROP\_MEMBERSHIP

To drop a membership, use:

```
struct ip_mreq mreq;
setsockopt(s, IPPROTO_IP, IP_DROP_MEMBERSHIP, &mreq, sizeof(mreq));
```

where *mreq* contains the same values as used to add the membership. Memberships are dropped when the socket is closed or the process exits.

**Returns:**

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

**Errors:**

## EADDRNOTAVAIL

You tried to create a socket with a network address for which no network interface exists.

## EISCONN

You tried to establish a connection on a socket that already has one or to send a datagram with the destination address specified, but the socket is already connected.

## ENOBUFS

The system ran out of memory for an internal data structure.

## ENOTCONN

You tried to send a datagram, but no destination address was specified and the socket hasn't been connected.



---

The following error specific to IP may occur when setting or getting IP options:

## EINVAL

An unknown socket option name was given. The IP option field was improperly formed — an option field was shorter than the minimum value or longer than the option buffer provided.

---

**See also:**

ICMP protocol

*connect()*, *getsockopt()*, *ioctl()*, *read()*, *recv()*, *recvfrom()*, *recvmsg()*, *send()*, *sendto()*, *setsockopt()*, *socket()*, *write()*

*RFC 791*

## Synopsis:

```
#include <sys/types.h>
#include <netinet/in.h>
#include <netinet6/ipsec.h>

int socket(PF_KEY,
 SOCK_RAW,
 PF_KEY_V2);
```

## Description:

IPsec is a security protocol for the Internet Protocol layer. It consists of these sub-protocols:

### AH (Authentication Header)

Guarantees the integrity of the IP packet and protects it from intermediate alteration or impersonation by attaching a cryptographic checksum computed by one-way hash functions.

### ESP (Encapsulated Security Payload)

Protects the IP payload from wire-tapping by encrypting it using secret-key cryptography algorithms.

IPsec has these modes of operation:

- Transport — protects peer-to-peer communication between end nodes.
- Tunnel — supports IP-in-IP encapsulation operation and is designed for security gateways, like VPN configurations.

## Kernel interface

The IPsec protocol behavior is controlled by these engines:

- Key management engine — accessed from an application using PF\_KEY sockets. The *RFC 2367* specification defines the PF\_KEY socket API.

- Policy engine — accessed with the PF\_KEY API, the *setsockopt()* operations, and the *sysctl()* interface. (The **sysctl** utility is a cover for the *sysctl()* function.) The *setsockopt()* function defines per-socket behavior and the *sysctl()* interface defines host-wide default behavior.

These engines are located in the socket manager. The socket manager implements the PF\_KEY interface and allows you to define IPsec policy similar to per-packet filters. Note that the socket manager code doesn't implement the dynamic encryption key exchange protocol IKE (Internet Key Exchange) — that implementation should be done at the application level (usually as daemons), using the previously described APIs.

### Policy management

The socket manager implements experimental policy management. You can manage the IPsec policy in these ways:

- Configure a per-socket policy using *setsockopt()*.
- Configure the socket manager packet filter-based policy using the PF\_KEY interface or via the **setkey** utility.

In this case, the default policy is allowed with the **setkey**. By configuring the policy to **default**, you can use the system-wide **sysctl** utility variables. (The **sysctl** utility displays various runtime options.)

If the socket manager finds no matching policy, the system-wide default value is applied.

For a list of *net.inet6.ipsec6.\** variables, see the **sysctl** utility in the *Utilities Reference*.

### Miscellaneous **sysctl** variables

The following variables are accessible via the **sysctl** utility for tweaking socket manager IPsec behavior:

| Name                                | Type    | Changeable? |
|-------------------------------------|---------|-------------|
| <i>net.inet.ipsec.ah_cleartos</i>   | Integer | Yes         |
| <i>net.inet.ipsec.ah_offsetmask</i> | Integer | Yes         |
| <i>net.inet.ipsec.dfbit</i>         | Integer | Yes         |
| <i>net.inet.ipsec.ecn</i>           | Integer | Yes         |
| <i>net.inet.ipsec.debug</i>         | Integer | Yes         |
| <i>net.inet6.ipsec6.ecn</i>         | Integer | Yes         |
| <i>net.inet6.ipsec6.debug</i>       | Integer | Yes         |

The variables are interpreted as follows:

*ipsec.ah\_cleartos*

When computing AH authentication data, the socket manager clears the type-of-service field in the IPv4 header if the value is set to a nonzero value. The variable tweaks AH behavior to interoperate with devices that implement *RFC 1826* AH. Set this to a nonzero value (clear the type-of-service field) if you want to conform to *RFC 2402*.

*ipsec.ah\_offsetmask*

When computing AH authentication data, the socket manager includes the 16-bit fragment offset field (including flag bits) in the IPv4 header, after computing a logical “AND” with the variable. This variable tweaks the AH behavior to interoperate with devices that implement *RFC 1826* AH. Set this value to zero (clear the fragment offset field during computation) if you want to conform to *RFC 2402*.

*ipsec.dfbit*

Configures the socket manager behavior for IPv4 IPsec tunnel encapsulation. The variable is supplied to conform to *RFC 2403* Chapter 6.1.



**If the value is set to: Then:**

|   |                                                              |
|---|--------------------------------------------------------------|
| 0 | The DF bit on the outer IPv4 header is cleared.              |
| 1 | The outer DF bit on the header is set from the inner DF bit. |
| 2 | The DF bit is copied from the inner header to the outer.     |

*ipsec.ecn* If set to nonzero, the IPv4 IPsec tunnel encapsulation/decapsulation behavior supports ECN (Explicit Congestion Notification), as documented in the IETF draft `draft-ietf-ipsec-ecn-02.txt`.

*ipsec.debug* If set to nonzero, debug messages are generated to the `syslog`.

Variables under the `net.inet6.ipsec6` tree have meaning similar to their `net.inet.ipsec` counterparts.

**Protocols**

Because the IPsec protocol works like a plugin to the INET and INET6 protocols, IPsec supports most of the protocols defined upon those IP-layer protocols. Some of the protocols, like ICMP or ICMP6, may behave differently with IPsec. This is because IPsec can prevent ICMP or ICMP6 routines from looking into the IP payload.

**Setting the policy**

You can set the policy manually by calling `setkey`, or set it permanently in `/etc/inetd.conf`. Valid policy settings include:

```
for setkey: -P direction discard
 -P direction ipsec request ...
 -P direction none
```

for `/etc/inetd.conf`:

```
direction bypass
direction entrust
direction ipsec request ...
```

where:

*direction* The direction in which the policy is applied. It's either **in** or **out**.

**bypass** (`/etc/inetd.conf` only) Bypass the IPsec processing and transmit the packet in clear text. This option is for privileged sockets.

**discard** (`setkey` only) Discard the packet matching indexes.

**entrust** (`/etc/inetd.conf` only) Consult the Security Policy Database (SPD) in the stack. The SPD is set by `setkey` (see the *Utilities Reference*).

*ipsec request ...*

Put the IPsec operation into the packet. You can specify one or more *request* strings using the following format:

```
protocol/mode/src-dst [/level]
```

For detailed descriptions of the arguments in the *request* string, see below.

**none** (`setkey` only) Don't put the IPsec operation into the packet.

#### Arguments for *request*

*protocol* One of:

- **ah** — Authentication Header. Guarantees the integrity of the IP packets and protects them from intermediate alteration or impersonation, by attaching cryptographic checksums computed by one-way hash functions.
- **esp** — Encapsulated Security Payload. Protects the IP payload from wire-tapping by encrypting it with secret key cryptography algorithms.
- **ipcomp** — IP Payload Compression Protocol.

*mode* Security protocol to be used, which is one of:

- **transport** — Protects peer-to-peer communication between end nodes.
- **tunnel** — Includes IP-in-IP encapsulation operations and is designed for security gateways, like VPN configurations.

*dst,*  
*src*

The “receiving node” (*dst*) and “sending node” (*src*) endpoint addresses of the Security Association (SA). When the direction specified is **in**, *dst* would represent this node and *src* the other node (peer).

If **transport** is specified as the *mode*, you can omit these values.

*level* One of:

- **default** — The stack should consult the system default policy that’s set by the **sysctl** utility.
- **require** — An SA is required whenever the kernel deals with the packet.
- **use** — Use an SA if it’s available; otherwise, keep the normal operation.
- **unique** — (**setkey** only) Similar to **require**, but adds the restriction that the SA for outbound traffic is used only for this policy.

You may need the identifier in order to relate the policy and the SA when you define the SA by manual keying. You can put the decimal number as the identifier after **unique**, such as:

**unique: number**

The value of *number* must be between 1 and 32767. If the request string is kept unambiguous, the *level* and slash prior to *level* can be omitted. However, you should specify them explicitly to avoid unintended behaviors.



---

If the *level* isn't specified in the **setkey** command, **unique** is used by default.

---

## Caveats:

The IPsec support is subject to change as the IPsec protocols develop.

There's no single standard for policy engine API, so the policy engine API described herein is just for KAME implementation.

The AH tunnel may not work as you might expect. If you configure the **require** policy against AH tunnel for inbound, tunneled packets will be rejected. This is because AH authenticates the encapsulating (outer) packet, not the encapsulated (inner) packet.

Under certain conditions, a truncated result may be returned from the socket manager from SADB\_DUMP and SADB\_SPDDUMP operations on a PF\_KEY socket. This occurs if there are too many database entries in the socket manager and the socket buffer for the PF\_KEY socket is too small. If you manipulate many IPsec key/policy database entries, increase the size of socket buffer.

## See also:

ICMP, ICMP6, INET6, IP, IPv6 protocols

*ioctl()*, *socket()*, *sysctl()*

*/etc/inetd.conf*, **setkey** in the *Utilities Reference*.

*RFC 2367, RFC 1826, RFC 2402, RFC 2403*

Detailed documentation about the IP security protocol may be found  
at the IPsec FAQ website at

<http://www.netbsd.org/Documentation/network/ipsec/>.

## ***ipsec\_dump\_policy()***

© 2004, QNX Software Systems Ltd.

*Generate a readable string from an IPsec policy specification*

### **Synopsis:**

```
#include <netinet6/ipsec.h>

char* ipsec_dump_policy(char *buf,
 char *delim);
```

### **Arguments:**

*buf*        A pointer to an IPsec policy structure **struct** **sadb\_x\_policy**.

*delim*     Delimiter string, usually a NULL which indicates a space (“ ”).

### **Library:**

**libipsec**

### **Description:**

The function *ipsec\_dump\_policy()* generates a readable string from an IPSEC policy specification. Please refer to *ipsec\_set\_policy()* for details about the policies.

The *ipsec\_dump\_policy()* function converts IPsec policy structure into a readable form. Therefore, *ipsec\_dump\_policy()* is the inverse of *ipsec\_set\_policy()*. If you set *delim* to NULL, a single whitespace is assumed. The function *ipsec\_dump\_policy()* returns a pointer to a dynamically allocated string. It is the caller's responsibility to reclaim the region, by using *free()*.

### **Returns:**

A pointer to dynamically allocated string, or NULL if an error occurs.

**Examples:**

See *ipsec\_set\_policy()*.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

IPsec, *ipsec\_get\_policylen()*, *ipsec\_set\_policy()*, *ipsec\_strerror()*  
**setkey** in the *Utilities Reference*

## ***ipsec\_get\_policylen()***

© 2004, QNX Software Systems Ltd.

*Get the length of the IPsec policy*

### **Synopsis:**

```
#include <netinet6/ipsec.h>

int ipsec_get_policylen(char *buf);
```

### **Arguments:**

*buf* A pointer to an IPsec policy structure **struct** **sadb\_x\_policy**.

### **Library:**

libipsec

### **Description:**

The function *ipsec\_get\_policylen()* gets the length of the IPsec policy. Please refer to *ipsec\_set\_policy()* for details about the policies.

You may want the length of the generated buffer when calling *setsockopt()*. The function *ipsec\_get\_policylen()* returns the length.

### **Returns:**

The size of the buffer, or a negative value if an error occurs.

### **Examples:**

See *ipsec\_set\_policy()*.

### **Classification:**

Unix

#### **Safety**

---

Cancellation point No

Interrupt handler Yes

*continued...*



**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

IPsec, *ipsec\_dump\_policy()*, *ipsec\_set\_policy()*, *ipsec\_strerror()*  
**setkey** in the *Utilities Reference*

## ***ipsec\_strerror()***

© 2004, QNX Software Systems Ltd.

*Error code for IPsec policy manipulation library*

### **Synopsis:**

```
#include <netinet6/ipsec.h>

const char *
ipsec_strerror(void);
```

### **Library:**

libipsec

### **Description:**

This *ipsec\_strerror()* function is used to obtain the error message string from the last failed ipsec call.

### **Returns:**

A pointer to an error message.



---

Don't modify the string that this function returns.

---

### **Examples:**

```
#include <netinet6/ipsec.h>
#include <sys/socket.h>
#include <stdio.h>
#include <malloc.h>
#include <string.h>

int
main(void)
{
 char *sadb;
 char *policy = "in discard";
 int len;

 sadb = ipsec_set_policy(policy, strlen(policy));

 if (sadb == NULL) {
 fprintf(stderr, "ipsec_set_policy: %s\n", ipsec_strerror());
 return 1;
 }
}
```

```
len = ipsec_get_policylen(sadb);
printf("len: %d\n", len);

policy = NULL;
policy = ipsec_dump_policy(sadb, NULL);

if (policy == NULL) {
 fprintf(stderr, "ipsec_dump_policy: %s\n", ipsec_strerror());
 return 1;
}

printf("policy: %s\n", policy);

free(policy);
free(sadb);

return 0;
}
```

## Classification:

Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

**IPsec**, *ipsec\_dump\_policy()*, *ipsec\_get\_policylen()*, *ipsec\_set\_policy()*, **setkey**

## ***ipsec\_set\_policy()***

© 2004, QNX Software Systems Ltd.

*Generate an IPsec policy specification structure from a readable string*

### **Synopsis:**

```
#include <netinet6/ipsec.h>

char* ipsec_set_policy(char *policy,
 int len);
```

### **Arguments:**

*len*           The length of the policy string.

*policy*        A string that describes a **struct sadb\_x\_policy** and optionally a **struct sadb\_x\_ipsecrequest**, formatted as described below.

### **Library:**

**libipsec**

### **Description:**

The function *ipsec\_set\_policy()* generates an IPsec policy specification structure, namely a **struct sadb\_x\_policy** and potentially a **struct sadb\_x\_ipsecrequest** from a human-readable policy specification. This function returns a pointer to the IPsec policy specification structure.



---

You should release the buffer returned by *ipsec\_set\_policy()* by calling *free()*. See the example below.

---

The policy is formatted as one of the following:

#### *direction discard*

The direction must be **in** or **out**. It specifies which direction the policy needs to be applied. With the discard policy, packets are dropped if they match the policy.

#### *direction entrust*

Consultation to SPD — defined by **setkey**.

*direction bypass*

Bypass the IPsec processing, i.e. packets are transmitted in clear. This is for privileged sockets.

*direction ipsec request ...*

The matching packets are subject to IPsec processing. The **ipsec** string can be followed by one or more request strings, which are formatted as below:

*protocol / mode / src - dst [/level]*

*protocol* Either **ah**, **esp**, or **ipcomp**.

*mode* Either **transport** or **tunnel**.

*src* and *dst* The IPsec endpoints; *src* is the sending node and *dst* is the receiving node. Therefore, when direction is **in**, *dst* is this node and *src* is the other node (peer).

*level* Either **default**, **use**, **require** or **unique**.

- **default** — the kernel should consult the system default policy defined by *sysctl()*.
- **use** — a relevant SA (security association) is used when available, since the kernel may perform IPsec operation against packets when possible. In this case, packets are transmitted in clear (when SA is not available), or encrypted (when SA is available).
- **require** — a relevant SA is required, since the kernel must perform IPsec operation against packets.
- **unique** is the same as **require**. However, it adds the restriction that the SA for outbound traffic is used only for this policy. You may need the identifier in order to relate the policy and the SA when you define the SA by manual keying. You put the decimal number as the identifier like:

**unique:** *number*

where *number* must be between 1 and 32767. If the request string is kept unambiguous, you can omit the *level* and the slash (“/”) prior to *level*. However, you should specify them explicitly to avoid unintended behavior. If *level* is omitted, it will be interpreted as default.

Here’s an example of policy information:

```
in discard
out ipsec esp/transport//require
in ipsec ah/transport//require
out ipsec esp/tunnel/10.1.1.2-10.1.1.1/use
in ipsec ipcom/transport//use esp/transport//use
```



---

It differs from the specification of **setkey**, where both **entrust** and **bypass** are not used. Please refer to **setkey** for detail.

---

## Returns:

A pointer to the allocated policy specification, or NULL if an error occurs.

## Examples:

```
#include <netinet6/ipsec.h>
#include <sys/socket.h>
#include <stdio.h>
#include <malloc.h>
#include <string.h>

int
main(void)
{
 char *sadb;
 char *policy = "in discard";
 int len;

 sadb = ipsec_set_policy(policy, strlen(policy));

 if (sadb == NULL) {
```

```
 fprintf(stderr, "ipsec_set_policy: %s\n", ipsec_strerror());
 return 1;
 }

 len = ipsec_get_policylen(sadb);
 printf("len: %d\n", len);

 policy = NULL;
 policy = ipsec_dump_policy(sadb, NULL);

 if (policy == NULL) {
 fprintf(stderr, "ipsec_dump_policy: %s\n", ipsec_strerror());
 return 1;
 }

 printf("policy: %s\n", policy);

 free(policy);
 free(sadb);

 return 0;
}
```

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

IPsec, *ipsec\_dump\_policy()*, *ipsec\_get\_policylen()*, *ipsec\_strerror()*  
**setkey** in the *Utilities Reference*

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(AF_INET6,
 SOCK_RAW,
 proto);
```

**Description:**

The IP6 protocol is the network-layer protocol used by the Internet Protocol version 6 family (AF\_INET6). Options may be set at the IP6 level when using higher-level protocols based on IP6 (such as TCP and UDP). It may also be accessed through a “raw socket” when developing new protocols, or special-purpose applications.

There are several IP6-level *setsockopt()/getsockopt()* options. They are separated into the basic IP6 sockets API (defined in *RFC2553*), and the advanced API (defined in *RFC2292*). The basic API looks very similar to the API presented in IP. The advanced API uses ancillary data and can handle more complex cases.



---

Specifying some of the socket options requires **root** privileges.

---

**Basic IP6 sockets API**

You can use the IPV6\_UNICAST\_HOPS option to set the hoplimit field in the IP6 header on unicast packets. If you specify -1, the socket manager uses the default value. If you specify a value of 0 to 255, the packet uses the specified value as its hoplimit. Other values are considered invalid and result in an error code of EINVAL. For example:

```
int hlim = 60; /* max = 255 */
setsockopt(s, IPPROTO_IPV6, IPV6_UNICAST_HOPS,
 &hlim, sizeof(hlim));
```



The IP6 multicasting is supported only on AF\_INET6 sockets of type SOCK\_DGRAM and SOCK\_RAW, and only on networks where the interface driver supports multicasting.

The IPV6\_MULTICAST\_HOPS option changes the hoplimit for outgoing multicast datagrams in order to control the scope of the multicasts:

```
unsigned int hlim; /* range: 0 to 255, default = 1 */
setsockopt(s, IPPROTO_IPV6, IPV6_MULTICAST_HOPS,
 &hlim, sizeof(hlim));
```

Datagrams with a hoplimit of 1 aren't forwarded beyond the local network. Multicast datagrams with a hoplimit of 0 won't be transmitted on any network, but may be delivered locally if the sending host belongs to the destination group and if multicast loopback hasn't been disabled on the sending socket (see below). Multicast datagrams with a hoplimit greater than 1 may be forwarded to other networks if a multicast router is attached to the local network.

For hosts with multiple interfaces, each multicast transmission is sent from the primary network interface. The IPV6\_MULTICAST\_IF option overrides the default for subsequent transmissions from a given socket:

```
unsigned int outif;
outif = if_nametoindex("ne0");
setsockopt(s, IPPROTO_IPV6, IPV6_MULTICAST_IF,
 &outif, sizeof(outif));
```

(The *outif* argument is an interface index of the desired interface, or 0 to specify the default interface.)

If a multicast datagram is sent to a group to which the sending host itself belongs (on the outgoing interface), a copy of the datagram is, by default, looped back by the IP6 layer for local delivery. The IPV6\_MULTICAST\_LOOP option gives the sender explicit control over whether or not subsequent datagrams are looped back:

```
u_char loop; /* 0 = disable, 1 = enable (default) */
setsockopt(s, IPPROTO_IPV6, IPV6_MULTICAST_LOOP,
 &loop, sizeof(loop));
```

This option improves performance for applications that may have no more than one instance on a single host (such as a router daemon), by eliminating the overhead of receiving their own transmissions. Don't use the `IPV6_MULTICAST_LOOP` option if there might be more than one instance of your application on a single host (e.g. a conferencing program), or if the sender doesn't belong to the destination group (e.g. a time-querying program).

A multicast datagram sent with an initial hoplimit greater than 1 may be delivered to the sending host on a different interface from that on which it was sent, if the host belongs to the destination group on that other interface. The loopback control option has no effect on such a delivery.

A host must become a member of a multicast group before it can receive datagrams sent to the group. To join a multicast group, use the `IPV6_JOIN_GROUP` option:

```
struct ipv6_mreq mreq6;
setsockopt(s, IPPROTO_IPV6, IPV6_JOIN_GROUP,
 &mreq6, sizeof(mreq6));
```

Note that the *mreq6* argument has the following structure:

```
struct ipv6_mreq {
 struct in6_addr ipv6mr_multiaddr;
 unsigned int ipv6mr_interface;
};
```

Set the *ipv6mr\_interface* member to 0 to choose the default multicast interface, or set it to the interface index of a particular multicast-capable interface if the host is multihomed. Membership is associated with a single interface; programs running on multihomed hosts may need to join the same group on more than one interface.

To drop a membership, use:

```
struct ipv6_mreq mreq6;
setsockopt(s, IPPROTO_IPV6, IPV6_LEAVE_GROUP,
 &mreq6, sizeof(mreq6));
```

The *mreq6* argument contains the same values as used to add the membership. Memberships are dropped when the socket is closed or the process exits.

The IPV6\_PORTRANGE option controls how ephemeral ports are allocated for SOCK\_STREAM and SOCK\_DGRAM sockets. For example:

```
int range = IPV6_PORTRANGE_LOW; /* see <netinet/in.h> */
setsockopt(s, IPPROTO_IPV6, IPV6_PORTRANGE, &range,
 sizeof(range));
```

The IPV6\_BINDV6ONLY option controls the behavior of the AF\_INET6 wildcard listening socket. The following example sets the option to 1:

```
int on = 1;
setsockopt(s, IPPROTO_IPV6, IPV6_BINDV6ONLY,
 &on, sizeof(on));
```

If you set the IPV6\_BINDV6ONLY option to 1, the AF\_INET6 wildcard listening socket accepts IP6 traffic only. If set to 0, the socket accepts IPv4 traffic as well, as if it were from an IPv4 mapped address, such as `::ffff:10.1.1.1`. Note that if you set the option to 0, IPv4 access control gets much more complicated. For example, even if you have no listening AF\_INET socket on port X, you'll end up accepting IPv4 traffic by an AF\_INET6 listening socket on the same port. The default value for this flag is copied at socket-instantiation time, from the *net.inet6.ip6.bindv6only* variable from the `sysctl` utility. The option affects TCP and UDP sockets only.

## Advanced IP6 sockets API

The advanced IP6 sockets API lets applications specify or obtain details about the IP6 header and extension headers on packets. The advanced API uses ancillary data for passing data to or from the socket manager.

There are also *setsockopt()* / *getsockopt()* options to get optional information on incoming packets:

- IPV6\_PKTINFO

- IPV6\_HOPLIMIT
- IPV6\_HOPOPTS
- IPV6\_DSTOPTS
- IPV6\_RTHDR

```
int on = 1;

setsockopt(fd, IPPROTO_IPV6, IPV6_PKTINFO,
 &on, sizeof(on));
setsockopt(fd, IPPROTO_IPV6, IPV6_HOPLIMIT,
 &on, sizeof(on));
setsockopt(fd, IPPROTO_IPV6, IPV6_HOPOPTS,
 &on, sizeof(on));
setsockopt(fd, IPPROTO_IPV6, IPV6_DSTOPTS,
 &on, sizeof(on));
setsockopt(fd, IPPROTO_IPV6, IPV6_RTHDR,
 &on, sizeof(on));
```

When any of these options are enabled, the corresponding data is returned as control information by *recvmsg()*, as one or more ancillary data objects.

If IPV6\_PKTINFO is enabled, the destination IP6 address and the arriving interface index are available via **struct in6\_pktinfo** on an ancillary data stream. You can pick the structure by checking for an ancillary data item by setting the *cmsg\_level* argument to IPPROTO\_IPV6 and the *cmsg\_type* argument to IPV6\_PKTINFO.

If IPV6\_HOPLIMIT is enabled, the hoplimit value on the packet is made available to the application. The ancillary data stream contains an integer data item with a *cmsg\_level* of IPPROTO\_IPV6 and a *cmsg\_type* of IPV6\_HOPLIMIT.

The *inet6\_option\_space()* family of functions help you parse ancillary data items for IPV6\_HOPOPTS and IPV6\_DSTOPTS. Similarly, the *inet6\_rthdr\_space()* family of functions help you parse ancillary data items for IPV6\_RTHDR.




---

The IPV6\_HOPOPTS and IPV6\_DSTOPTS values may appear multiple times on an ancillary data stream (note that the behavior is slightly different from the specification). Other ancillary data items appear no more than once.

---

You can pass ancillary data items with normal payload data, using the *sendmsg()* function. Ancillary data items are parsed by the socket manager, and are used to construct the IP6 header and extension headers. For the *msg\_level* values listed above, the ancillary data format is the same as the inbound case.

Additionally, you can specify a IPV6\_NEXTHOP data object. The IPV6\_NEXTHOP ancillary data object specifies the next hop for the datagram as a socket address structure. In the *cmsghdr* structure containing this ancillary data, the *msg\_level* argument is IPPROTO\_IPV6, the *msg\_type* argument is IPV6\_NEXTHOP, and the first byte of *msg\_data* is the first byte of the socket address structure.

If the socket address structure contains an IP6 address (e.g. the *sin6\_family* argument is AF\_INET6 ), then the node identified by that address must be a neighbor of the sending host. If that address equals the destination IP6 address of the datagram, then this is equivalent to the existing SO\_DONTROUTE socket option.

For applications that don't, or can't use the *sendmsg()* or the *recvmsg()* function, the IPV6\_PKTOPTIONS socket option is defined. Setting the socket option specifies any of the optional output fields:

```
setsockopt(fd, IPPROTO_IPV6, IPV6_PKTOPTIONS,
 &buf, len);
```

The *buf* argument points to a buffer containing one or more ancillary data objects; the *len* argument is the total length of all these objects. The application fills in this buffer exactly as if the buffer were being passed to the *sendmsg()* function as control information.

The options set by calling *setsockopt()* for IPV6\_PKTOPTIONS are called “sticky” options because once set, they apply to all packets sent on that socket. The application can call *setsockopt()* again to change

all the sticky options, or it can call *setsockopt()* with a length of 0 to remove all the sticky options for the socket.

The corresponding receive option:

```
getsockopt(fd, IPPROTO_IPV6, IPV6_PKTOPTIONS,
 &buf, &len);
```

returns a buffer with one or more ancillary data objects for all the optional receive information that the application has previously specified that it wants to receive. The *buf* argument points to the buffer that the call fills in. The *len* argument is a pointer to a value-result integer; when the function is called, the integer specifies the size of the buffer pointed to by *buf*, and on return this integer contains the actual number of bytes that were stored in the buffer. The application processes this buffer exactly as if it were returned by *recvmsg()* as control information.

### Advanced API and TCP sockets

When using *getsockopt()* with the `IPV6_PKTOPTIONS` option and a TCP socket, only the options from the most recently received segment are retained and returned to the caller, and only after the socket option has been set. The application isn't allowed to specify ancillary data in a call to *sendmsg()* on a TCP socket, and none of the ancillary data described above is ever returned as control information by *recvmsg()* on a TCP socket.

### Conflict resolution

In some cases, there are multiple APIs defined for manipulating an IP6 header field. A good example is the outgoing interface for multicast datagrams: it can be manipulated by `IPV6_MULTICAST_IF` in the basic API, by `IPV6_PKTINFO` in the advanced API, and by the *sin6\_scope\_id* field of the socket address structure passed to the *sendto()* function.

In QNX Neutrino, when conflicting options are given to the socket manager, the socket manager gets the value in the following order:

- 1 options specified by using ancillary data

- 2 options specified by a sticky option of the advanced API
- 3 options specified by using the basic API
- 4 options specified by a socket address.



---

The conflict resolution is undefined in the API specification and depends on the implementation.

---

### Raw IP6 Sockets

Raw IP6 sockets are connectionless, and are normally used with *sendto()* and *recvfrom()*, although you can also use *connect()* to fix the destination for future packets (in which case you can use *read()* or *recv()*, and *write()* or *send()*).

If *proto* is 0, the default protocol IPPROTO\_RAW is used for outgoing packets, and only incoming packets destined for that protocol are received. If *proto* is nonzero, that protocol number is used on outgoing packets and to filter incoming packets.

Outgoing packets automatically have an IP6 header prepended to them (based on the destination address and the protocol number the socket is created with). Incoming packets are received without the IP6 header or extension headers.

All data sent via raw sockets *must* be in network byte order; all data received via raw sockets is in network-byte order. This differs from the IPv4 raw sockets, which didn't specify a byte ordering and typically used the host's byte order.

Another difference from IPv4 raw sockets is that complete packets (i.e. IP6 packets with extension headers) can't be read or written using the IP6 raw sockets API. Instead, ancillary data objects are used to transfer the extension headers, as described above.

All fields in the IP6 header that an application might want to change (i.e. everything other than the version number) can be modified using ancillary data and/or socket options by the application for output. All fields in a received IP6 header (other than the version number and

Next Header fields) and all extension headers are also made available to the application as ancillary data on input. Hence, there's no need for a socket option similar to the IPv4 `IP_HDRINCL` socket option.

When writing to a raw socket, the socket manager automatically fragments the packet if the size exceeds the path MTU, inserting the required fragmentation headers. On input, the socket manager reassembles received fragments, so the reader of a raw socket never sees any fragment headers.

Most IPv4 implementations give special treatment to a raw socket created with a third argument to `socket()` of `IPPROTO_RAW`, whose value is normally 255. We note that this value has no special meaning to an IP6 raw socket (and the IANA currently reserves the value of 255 when used as a next-header field).

For ICMP6 raw sockets, the socket manager calculates and inserts the mandatory ICMP6 checksum.

For other raw IP6 sockets (i.e. for raw IP6 sockets created with a third argument other than `IPPROTO_ICMPV6`), the application must:

- 1 Set the new `IPV6_CHECKSUM` socket option to have the socket manager compute and store a pseudo header checksum for output.
- 2 Verify the received pseudo header checksum on input, discarding the packet if the checksum is in error.

This option prevents applications from having to perform source-address selection on the packets they send. The checksum incorporates the IP6 pseudo-header, defined in Section 8.1 of *RFC 2460*. This new socket option also specifies an integer offset into the user data of where the checksum is located.

```
int offset = 2;
setsockopt(fd, IPPROTO_IPV6, IPV6_CHECKSUM,
 &offset, sizeof(offset));
```

By default, this socket option is disabled. Setting the offset to -1 also disables the option. Disabled means:



- 1 The socket manager won't calculate and store a checksum for outgoing packets.
- 2 The socket manager kernel won't verify a checksum for received packets.



- Since the checksum is always calculated by the socket manager for an ICMP6 socket, applications can't generate ICMPv6 packets with incorrect checksums (presumably for testing purposes) using this API.
  - The IPV6\_NEXTHOP object/option isn't fully implemented.
- 

### See also:

*getsockopt()*, ICMP6 protocol, INET6 protocol, *recv()*, *send()*, *setsockopt()*

## ***isalnum()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's alphanumeric*

---

### **Synopsis:**

```
#include <ctype.h>

int isalnum(int c);
```

### **Arguments:**

*c* The character you want to test.

### **Library:**

`libc`

### **Description:**

The *isalnum()* function tests if the argument *c* is an alphanumeric character (**a** to **z**, **A** to **Z**, or **0** to **9**). An alphanumeric character is any character for which *isalpha()* or *isdigit()* is true.

### **Returns:**

Nonzero if *c* is a letter or decimal digit; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

int main(void)
{
 if(isalnum(getchar())) {
 printf("That's alpha-numeric!\n");
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The result is only valid for `char` arguments and EOF.

## See also:

*isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isalpha()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's alphabetic*

---

### **Synopsis:**

```
#include <ctype.h>

int isalpha(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

`libc`

### **Description:**

The *isalpha()* function tests if the argument *c* is an alphabetic character (**a** to **z** and **A** to **Z**). An alphabetic character is any character for which *isupper()* or *islower()* is true.

### **Returns:**

Nonzero if *c* is an alphabetic character; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

int main(void)
{
 if(isalpha(getchar())) {
 printf("That's alphabetic\n");
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum(), iscntrl(), isdigit(), isgraph(), islower(), isprint(), ispunct(), isspace(), isupper(), isxdigit(), tolower(), toupper()*

## ***isascii()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's a 7-bit ASCII character*

### **Synopsis:**

```
#include <ctype.h>

int isascii(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

libc

### **Description:**

The *isascii()* function tests for an ASCII character (in the range 0 to 127).

### **Returns:**

Nonzero if *c* is an ASCII character; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 0x80, 'Z' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isascii(the_chars[i])) {
 printf("Char %c is an ASCII character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not an ASCII character\n",
 the_chars[i]);
 }
 }
}
```

```
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is an ASCII character
Char is not an ASCII character
Char Z is an ASCII character
```

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalpha()*, *isalnum()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isatty()***

© 2004, QNX Software Systems Ltd.

*Test to see if a file descriptor is associated with a terminal*

### **Synopsis:**

```
#include <unistd.h>

int isatty(int fdes);
```

### **Arguments:**

*fdes*     The file descriptor that you want to test.

### **Library:**

libc

### **Description:**

The *isatty()* function allows the calling process to determine if the file descriptor *fdes* is associated with a terminal.

### **Returns:**

- 0     The *fdes* file descriptor doesn't refer to a terminal.
- 1     The *fdes* file descriptor refers to a terminal.

### **Examples:**

```
/*
 * The following program exits with a status of
 * EXIT_SUCCESS if stderr is a tty; otherwise,
 * EXIT_FAILURE
 */
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(void)
{
 return(isatty(3) ? EXIT_SUCCESS : EXIT_FAILURE);
}
```



## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*open()*

## ***iscntrl()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's a control character*

### **Synopsis:**

```
#include <ctype.h>

int iscntrl(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

libc

### **Description:**

The *iscntrl()* function tests for any control character. An ASCII control character is any character whose value is between 0 and 31.

### **Returns:**

Nonzero if *c* is a control character; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 0x09, 'Z' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(iscntrl(the_chars[i])) {
 printf("Char %c is a Control character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a Control character\n",
 the_chars[i]);
 }
 }
}
```

```
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is not a Control character
Char is a Control character
Char Z is not a Control character
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*,  
*isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isdigit()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's a decimal digit*

### **Synopsis:**

```
#include <ctype.h>

int isdigit(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

libc

### **Description:**

The *isdigit()* function tests for a decimal digit (characters 0 through 9).

### **Returns:**

Nonzero if *c* is a decimal digit; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', '5', '$' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isdigit(the_chars[i])) {
 printf("Char %c is a digit character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a digit character\n",
 the_chars[i]);
 }
 }
}
```

```
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is not a digit character
Char 5 is a digit character
Char $ is not a digit character
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isfdtype()***

© 2004, QNX Software Systems Ltd.

*Determine whether a file descriptor refers to a socket*

### **Synopsis:**

```
#include <sys/stat.h>

int isfdtype(int filedes,
 int fdtype);
```

### **Arguments:**

- filedes*    The file descriptor that you want to test.
- fdtype*    The properties you want to test for. The valid values for *fdtype* include:
- S\_IFSOCK — test whether *filedes* is a socket.

### **Library:**

libc

### **Description:**

The *isfdtype()* function determines whether the file descriptor *filedes* has the properties identified by *fdtype*.

### **Returns:**

- 1        The *filedes* file descriptor matches *fdtype*.
- 0        The *filedes* file descriptor doesn't match *fdtype*.
- 1       An error occurred (*errno* is set).

### **Errors:**

EBADF    Invalid file descriptor *filedes*.

## Classification:

POSIX 1003.1g (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isatty()*, *socket()*, *stat()*

## ***isgraph()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's any printable character except a space*

### **Synopsis:**

```
#include <ctype.h>

int isgraph(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

`libc`

### **Description:**

The *isgraph()* function tests for any printable character except a space (' '). The *isprint()* function is similar, except that the space character is also included.

### **Returns:**

Nonzero if *c* is a printable character (except a space); otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 0x09, ' ', 0x7d };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isgraph(the_chars[i])) {
 printf("Char %c is a printable character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a printable character\n",
 the_chars[i]);
 }
 }
}
```



```
 the_chars[i]);
 }
}
return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is a printable character
Char is not a printable character
Char is not a printable character
Char } is a printable character
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isinf()*, *isinff()***

© 2004, QNX Software Systems Ltd.

*Test for infinity*

### **Synopsis:**

```
#include <math.h>

int isinf (double x);

int isinff (float x);
```

### **Arguments:**

*x*     The number that you want to test.

### **Library:**

libm

### **Description:**

The *isinf()* and *isinff()* functions test to see if a number is “infinity.”

### **Returns:**

1       The value of *x* is infinity.  
≠ 1     The value of *x* isn't infinity.

### **Examples:**

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b, c, d;

 a = 2;
 b = -0.5;
 c = NAN;
 fp_exception_mask(_FP_EXC_DIVZERO, 1);
 d = 1.0/0.0;
 printf("%f is %s \n", a,
```

```
 (isinf(a) ? "infinite" : "not infinite");
printf("%f is %s \n", b,
 (isinf(b) ? "infinite" : "not infinite");
printf("%f is %s \n", c,
 (isinf(c) ? "infinite" : "not infinite");
printf("%f is %s \n", d,
 (isinf(d) ? "infinite" : "not infinite");

return(0);
}
```

produces the output:

```
2.000000 is not infinite
-0.500000 is not infinite
NaN is not infinite
Inf is infinite
```

## Classification:

Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*sin()*, *sinh()*

## ***islower()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's a lowercase letter*

### **Synopsis:**

```
#include <ctype.h>

int islower(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

`libc`

### **Description:**

The *islower()* function tests for any lowercase letter **a** through **z**.

### **Returns:**

Nonzero if *c* is a lowercase letter; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 'a', 'z', 'Z' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(islower(the_chars[i])) {
 printf("Char %c is a lowercase character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a lowercase character\n",
 the_chars[i]);
 }
 }
}
```

```
 return EXIT_SUCCESS;
 }
```

produces the output:

```
Char A is not a lowercase character
Char a is a lowercase character
Char z is a lowercase character
Char Z is not a lowercase character
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isnan()*, *isnanf()***

© 2004, QNX Software Systems Ltd.

*Test for not-a-number (NaN)*

---

### **Synopsis:**

```
#include <math.h>

int isnan (double x);

int isnanf (float x);
```

### **Arguments:**

*x*     The number you want to test.

### **Library:**

libm

### **Description:**

The *isnan()* and *isnanf()* functions determine if *x* is Not-A-Number (NaN).

### **Returns:**

1       The value of *x* is NaN.  
≠ 1     The value of *x* is a number.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b, c, d;

 a = 2;
 b = -0.5;
 c = NAN;
```

```
fp_exception_mask(_FP_EXC_DIVZERO, 1);
d = 1.0/0.0;
printf("%f is %s \n", a,
 (isnan(a)) ? "not a number" : "a number");
printf("%f is %s \n", b,
 (isnan(b)) ? "not a number" : "a number");
printf("%f is %s \n", c,
 (isnan(c)) ? "not a number" : "a number");
printf("%f is %s \n", d,
 (isnan(d)) ? "not a number" : "a number");
return EXIT_SUCCESS;
}
```

produces the output:

```
2.000000 is a number
-0.500000 is a number
NAN is not a number
Inf is a number
```

## Classification:

*isnan()* is standard Unix; *isnanf()* is ANSI (draft)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*finite()*

## ***isprint()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's any printable character, including a space*

### **Synopsis:**

```
#include <ctype.h>

int isprint(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

`libc`

### **Description:**

The *isprint()* function tests for any printable character, including a space ( ' ' ). The *isgraph()* function is similar, except that the space character is excluded from the character set being tested.

### **Returns:**

Nonzero if *c* is a printable character; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 0x09, ' ', 0x7d };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isprint(the_chars[i])) {
 printf("Char %c is a printable character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a printable character\n",
 the_chars[i]);
 }
 }
}
```



```
 the_chars[i]);
 }
}
return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is a printable character
Char is not a printable character
Char is a printable character
Char } is a printable character
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***ispunct()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's any punctuation character*

### **Synopsis:**

```
#include <ctype.h>

int ispunct(int c);
```

### **Arguments:**

*c* The character you want to test.

### **Library:**

libc

### **Description:**

The *ispunct()* function tests for any punctuation character such as a comma (,) or a period (.).

### **Returns:**

Nonzero if *c* is punctuation; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', '!', '.', ',', ':', ';' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(ispunct(the_chars[i])) {
 printf("Char %c is a punctuation character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a punctuation character\n",
 the_chars[i]);
 }
 }
}
```

```
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is not a punctuation character
Char ! is a punctuation character
Char . is a punctuation character
Char , is a punctuation character
Char : is a punctuation character
Char ; is a punctuation character
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *isspace()*, *isupper()*, *isxdigit()*, *tolower()*, *toupper()*

## ***isspace()***

© 2004, QNX Software Systems Ltd.

*Test a character to see if it's a whitespace character*

### **Synopsis:**

```
#include <ctype.h>

int isspace(int c);
```

### **Arguments:**

*c*     The character you want to test.

### **Library:**

`libc`

### **Description:**

The *isspace()* function tests for the following whitespace characters:

|      |                     |
|------|---------------------|
| ' '  | space               |
| '\f' | form feed           |
| '\n' | newline or linefeed |
| '\r' | carriage return     |
| '\t' | horizontal tab      |
| '\v' | vertical tab        |

### **Returns:**

Nonzero if *c* is a whitespace character; otherwise, zero.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', 0x09, ' ', 0x7d };
```

```
#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isspace(the_chars[i])) {
 printf("Char %c is a space character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not a space character\n",
 the_chars[i]);
 }
 }

 return EXIT_SUCCESS;
}
```

This program produces the output:

```
Char A is not a space character
Char is a space character
Char is a space character
Char } is not a space character
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*isalnum(), isalpha(), iscntrl(), isdigit(), isgraph(), islower(), isprint(),  
ispunct(), isupper(), isxdigit(), tolower(), toupper()*

## Synopsis:

```
#include <ctype.h>

int isupper(int c);
```

## Arguments:

*c* The character you want to test.

## Library:

libc

## Description:

The *isupper()* function tests for any uppercase letter **A** through **Z**.

## Returns:

Nonzero if *c* is an uppercase letter; otherwise, zero.

## Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <ctype.h>

char the_chars[] = { 'A', 'a', 'z', 'Z' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isupper(the_chars[i])) {
 printf("Char %c is an uppercase character\n",
 the_chars[i]);
 } else {
 printf("Char %c is not an uppercase character\n",
 the_chars[i]);
 }
 }
}
```

```
 return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is an uppercase character
Char a is not an uppercase character
Char z is not an uppercase character
Char Z is an uppercase character
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isxdigit()*, *tolower()*, *toupper()*



**Synopsis:**

```
#include <wctype.h>

int iswalnum(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswalnum()* function tests if the argument *wc* is an alphanumeric wide character of the class `alpha` or `digit`. In the C locale, they're `a` to `z`, `A` to `Z`, `0` to `9`.

**Returns:**

A nonzero value if the character is a member of the class `alpha` or `digit`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswalpha(wint_t wc);
```

**Arguments:**

*wc* The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswalpha()* function tests if the argument *wc* is an alphabetic wide character of the class `a1pha`. In the C locale, they are: `a` to `z`, `A` to `Z`.

**Returns:**

A nonzero value if the character is a member of the class `a1pha`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswcntrl(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswcntrl()* function tests if the argument *wc* is a control wide character of the class `cntrl`. In the C locale, this class consists of the ASCII characters from 0 through 31.

**Returns:**

A nonzero value if the character is a member of the class `cntrl`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswctype(wint_t wc,
 wctype_t charclass);
```

**Arguments:**

*wc*            The wide character you want to test.

*charclass*    The character class you want to test for. Get this class by calling *wctype()*.

**Library:**

`libc`

**Description:**

The *iswctype()* function tests if the argument *wc* is a member of one or several character classes.

| <b>This function:</b>       | <b>Is equivalent to:</b>                        |
|-----------------------------|-------------------------------------------------|
| <code>iswalnum( wc )</code> | <code>iswctype( wc , wctype( "alnum" ) )</code> |
| <code>iswalpha( wc )</code> | <code>iswctype( wc , wctype( "alpha" ) )</code> |
| <code>ispunct( wc )</code>  | <code>iswctype( wc , wctype( "punct" ) )</code> |



The results are unreliable if you didn't use *wctype()* to obtain *charclass*, or if a call to *setlocale()* affects LC\_CTYPE.

---

**Returns:**

A nonzero value if the character is a member of the specified character class (or classes), or zero if the character isn't a member or *charclass* is 0.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.



**Synopsis:**

```
#include <wctype.h>

int iswdigit(wint_t wc);
```

**Arguments:**

*wc*     The wide character that you want to test.

**Library:**

`libc`

**Description:**

The *iswdigit()* function tests if the argument *wc* is a decimal digit wide character of the class `digit`. In the C locale, this class consists of the characters 0 through 9.

**Returns:**

A nonzero value if the character is a member of the class `digit`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

*Test a wide character to see if it's any printable character except space*

**Synopsis:**

```
#include <wctype.h>

int iswgraph(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswgraph()* function tests if the argument *wc* is a graphical wide character of the class **graph**. In the C locale, this class consists of all the printable characters, except the space character.

**Returns:**

A nonzero value if the character is a member of the class **graph**, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswlower(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswlower()* function tests if the argument *wc* is a lowercase wide character of the class `lower`. In the C locale, this class consists of the characters from `a` through `z`.

**Returns:**

A nonzero value if the character is a member of the class `lower`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

*Test a wide character to see if it's any printable character, including space*

**Synopsis:**

```
#include <wctype.h>

int iswprint(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswprint()* function tests if the argument *wc* is a printable wide character of the class `print`. In the C locale, this class consists of all the printable characters, including the space character.

**Returns:**

A nonzero value if the character is a member of the class `print`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.



**Synopsis:**

```
#include <wctype.h>

int iswpunct(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswpunct()* function tests if the argument *wc* is a punctuation wide character of the class `punct`. In the C locale, this class includes the comma (,) and the period (.), among others.

**Returns:**

A nonzero value if the character is a member of the class `punct`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswspace(wint_t wc);
```

**Arguments:**

*wc* The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswspace()* function tests if the argument *wc* is a whitespace wide character of the class **space**. In the C locale, this class includes the space character, `\f` (form feed), `\n` (newline or linefeed), `\r` (carriage return), `\t` (horizontal tab), and `\v` (vertical tab).

**Returns:**

A nonzero value if the character is a member of the class **space**, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswupper(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswupper()* function tests if the argument *wc* is an uppercase wide character of the class **upper**. In the C locale, this class includes the characters from **A** through **Z**.

**Returns:**

A nonzero value if the character is a member of the class **upper**, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wctype.h>

int iswxdigit(wint_t wc);
```

**Arguments:**

*wc*     The wide character you want to test.

**Library:**

`libc`

**Description:**

The *iswxdigit()* function tests if the argument *wc* is a hexadecimal wide character of the class `xdigit`. In the C locale, this class includes the characters 0 to 9, and A to F.

**Returns:**

A nonzero value if the character is a member of the class `xdigit`, or 0 otherwise.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The result is valid only for `wchar_t` arguments and WEOF.

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.



## Synopsis:

```
#include <ctype.h>

int isxdigit(int c);
```

## Arguments:

*c* The character you want to test.

## Library:

libc

## Description:

The *isxdigit()* function tests for any hexadecimal-digit character. These characters are the digits 0 through 9 and the letters **a** through **f** (or **A** through **F**).

## Returns:

Nonzero if *c* is a hexadecimal digit; otherwise, zero.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char the_chars[] = { 'A', '5', '$' };

#define SIZE sizeof(the_chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
 if(isxdigit(the_chars[i])) {
 printf("Char %c is a hexadecimal digit",
 the_chars[i]);
 } else {
 printf("Char %c is not a hexadecimal digit",
```

```
 the_chars[i]);
 }
}
return EXIT_SUCCESS;
}
```

produces the output:

```
Char A is a hexadecimal digit character
Char 5 is a hexadecimal digit character
Char $ is not a hexadecimal digit character
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *isctrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *tolower()*, *toupper()*

## Synopsis:

```
#include <stdlib.h>

char* itoa(int value,
 char* buffer,
 int radix);
```

## Arguments:

*value*     The value to convert into a string.

*buffer*    A buffer in which the function stores the string. The size of the buffer must be at least:

$8 \times \text{sizeof}( \text{int} ) + 1$

bytes when converting values in base 2 (binary).

*radix*     The base to use when converting the number.

If the value of *radix* is 10, and *value* is negative, then a minus sign is prepended to the result.

## Library:

`libc`

## Description:

The *itoa()* function converts the integer *value* into the equivalent string in base *radix* notation, storing the result in the specified *buffer*. The function terminates the string with a NUL character.

## Returns:

A pointer to the resulting string.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char buffer[20];
 int base;

 for(base = 2; base <= 16; base += 2) {
 printf("%2d %s\n", base,
 itoa(12765, buffer, base));
 }

 return EXIT_SUCCESS;
}
```

produces the output:

```
 2 11000111011101
 4 3013131
 6 135033
 8 30735
10 12765
12 7479
14 491b
16 31dd
```

**Classification:**

QNX 4

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atoi(), atol(), ltoa(), sscanf(), strtol(), strtoul(), ultoa(), utoa()*

## ***j0(), j0f()***

© 2004, QNX Software Systems Ltd.

*Compute a Bessel function of the first kind*

### **Synopsis:**

```
#include <math.h>

double j0(double x);

float j0f(float x);
```

### **Arguments:**

*x* The number that you want to compute the Bessel function for.

### **Library:**

```
libbessel
```

### **Description:**

Compute the Bessel function of the first kind for *x*.

### **Returns:**

The result of the Bessel function of *x*.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdlib.h>
#include <stdio.h>
#include <math.h>

int main(void)
{
 double x, y, z;

 x = j0(2.4);
 y = y1(1.58);
 z = jn(3, 2.4);
```

```
printf("j0(2.4) = %f, y1(1.58) = %f\n", x, y);
printf("jn(3,2.4) = %f\n", z);

return EXIT_SUCCESS;
}
```

## Classification:

*j0()* is standard Unix; *j0f()* is ANSI (draft)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno, j1(), jn(), y0(), y1(), yn()*

## ***j1(), j1f()***

© 2004, QNX Software Systems Ltd.

*Compute a Bessel function of the first kind*

### **Synopsis:**

```
#include <math.h>

double j1(double x);

float j1f(float x);
```

### **Arguments:**

*x* The number that you want to compute the Bessel function for.

### **Library:**

`libbessel`

### **Description:**

Compute the Bessel function of the first kind for *x*.

### **Returns:**

The result of the Bessel function of *x*.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Classification:**

*j1()* is standard Unix; *j1f()* is ANSI (draft)

#### **Safety**

---

Cancellation point No

Interrupt handler No

*continued...*



**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*errno, j0(), jn(), y0(), y1(), yn()*

## ***$j_n()$ , $j_n f()$***

© 2004, QNX Software Systems Ltd.

*Compute a Bessel function of the first kind*

---

### **Synopsis:**

```
#include <math.h>

double jn(int n, double x);

float jnf(int n, float x);
```

### **Arguments:**

*n, x*    The numbers that you want to compute the Bessel function for.

### **Library:**

`libbessel`

### **Description:**

Compute the Bessel function of the first kind for *n* and *x*.

### **Returns:**

The result of the Bessel function of *n* and *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdlib.h>
#include <stdio.h>
#include <math.h>

int main(void)
{
 double x, y, z;

 x = j0(2.4);
```

```
 y = y1(1.58);
 z = jn(3, 2.4);

 printf("j0(2.4) = %f, y1(1.58) = %f\n", x, y);
 printf("jn(3,2.4) = %f\n", z);

 return EXIT_SUCCESS;
}
```

## Classification:

*jn()* is standard Unix; *jnf()* is ANSI (draft)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *j0()*, *j1()*, *y0()*, *y1()*, *yn()*

## ***jrand48()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random signed long integer in a thread-safe manner*

### **Synopsis:**

```
#include <stdlib.h>

long jrand48(unsigned short xsubi[3]);
```

### **Arguments:**

*xsubi* An array that comprises the 48 bits of the initial value that you want to use.

### **Library:**

libc

### **Description:**

The *jrand48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a signed **long** integer uniformly distributed over the interval  $[-2^{31}, 2^{31})$ . It's a thread-safe version of *mrnd48()*.

The *xsubi* array should contain the desired initial value; this makes *jrand48()* thread-safe, and lets you start a sequence of random numbers at any known value.

### **Returns:**

A pseudo-random **long** integer.

### **Classification:**

Standard Unix

#### **Safety**

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*drand48(), erand48(), lcong48(), lrand48(), mrand48(), nrand48(), seed48(), srand48()*

## ***kill()***

© 2004, QNX Software Systems Ltd.

*Send a signal to a process or a group of processes*

### **Synopsis:**

```
#include <sys/types.h>
#include <signal.h>

int kill(pid_t pid,
 int sig);
```

### **Arguments:**

*pid* The ID of the process or process group that you want to send a signal to:

#### **If *pid* is: Then *sig* is sent to:**

---

|     |                                                                         |
|-----|-------------------------------------------------------------------------|
| > 0 | The single process with that process ID                                 |
| 0   | All processes that are in the same process group as the sending process |
| < 0 | Every process that's a member of the process group <i>-pid</i>          |

*sig* Zero, or the signal that you want to send. For a complete list of signals, see “POSIX signals” in the documentation for *SignalAction()*.

### **Library:**

`libc`

### **Description:**

The *kill()* function sends the signal *sig* to a process or group of processes specified by *pid*. If *sig* is zero, no signal is sent, but the *pid* is still checked for validity.

For a process to have permission to send a signal to a process, the real or effective user ID of the sending process must either:

- match the real or effective user ID of the receiving process

Or:

- equal zero.

If the value of *pid* causes *sig* to be generated for the sending process, and if *sig* isn't blocked, either *sig* or at least one pending unblocked signal is delivered before the *kill* function returns.

This call doesn't block. However, in the network case, lower priority threads may run.

## Returns:

Zero, or -1 if an error occurs (*errno* is set).

## Errors:

|        |                                                                                   |
|--------|-----------------------------------------------------------------------------------|
| EAGAIN | Insufficient system resources are available to deliver the signal.                |
| EINVAL | The <i>sig</i> is invalid.                                                        |
| EPERM  | The process doesn't have permission to send this signal to any receiving process. |
| ESRCH  | The given <i>pid</i> doesn't exist.                                               |

## Examples:

See *sigprocmask()*.

## Classification:

POSIX 1003.1

### Safety

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*getpid(), killpg(), setsid(), sigaction(), signal(), SignalKill(), sigqueue()*



## Synopsis:

```
#include <sys/types.h>
#include <signal.h>

int killpg(pid_t pgrp,
 int sig);
```

## Arguments:

*pid*     The ID of the process group that you want to send a signal to.

*sig*     Zero, or the signal that you want to send. For a complete list of signals, see “POSIX signals” in the documentation for *SignalAction()*.

## Library:

libc

## Description:

The *killpg()* function sends the signal *sig* to the process group specified by *pgrp*. If *sig* is zero, no signal is sent, but *pgrp* is still checked for validity.

If *pgrp* is greater than 1, *killpg (pgrp, sig)* is equivalent to *kill (-pgrp, sig)*.

## Returns:

Zero, or -1 if an error occurs (*errno* is set).

## Errors:

EAGAIN     Insufficient system resources are available to deliver the signal.

EINVAL     The signal *sig* is invalid or not supported.

- EPERM      The process doesn't have permission to send this signal to any receiving process.
- ESRCH      No process group can be found corresponding to the specified *pgrp* or *pgrp* is less than or equal to 1.

**Examples:**

See *sigprocmask()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*getpid()*, *kill()*, *setsid()*, *sigaction()*, *signal()*, *SignalKill()*, *sigqueue()*

**Synopsis:**

```
#include <stdlib.h>

long labs(long j);
```

**Arguments:**

*j*      The number you want the absolute value of.

**Library:**

libc

**Description:**

The *labs()* function returns the absolute value of its long-integer argument *j*.

**Returns:**

The absolute value of *j*.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 long x, y;

 x = -50000;
 y = labs(x);
 printf("labs(%d) = %d\n", x, y);
 return EXIT_SUCCESS;
}
```

produces the output:

```
labs(-50000) = 50000
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*abs()*, *cabs()*, *fabs()*

### **Synopsis:**

```
#include <sys/types.h>
#include <unistd.h>

int lchown(const char * path,
 uid_t owner,
 gid_t group);
```

### **Arguments:**

*path*        The name of the file whose ownership you want to change.

*owner*       The user ID of the new owner.

*group*       The group ID of the new owner.

### **Library:**

`libc`

### **Description:**

The *lchown()* function changes the user ID and group ID of the file specified by *path* to be the numeric values contained in *owner* and *group*, respectively. It's similar to *chown()*, except in the case where the named file is a symbolic link. In this case, *lchown()* changes the ownership of the symbolic link file itself, while *chown()* changes the ownership of the file or directory to which the symbolic link refers.

Only processes with an effective user ID equal to the user ID of the file or with appropriate privileges (for example, the superuser) may change the ownership of a file.

In QNX Neutrino, the `_POSIX_CHOWN_RESTRICTED` flag is enforced. This means that only the superuser may change the ownership of a file. The group of a file may be changed by the superuser, or also by a process with the effective user ID equal to the user ID of the file, if (and only if) *owner* is equal to the user ID of the file and *group* is equal to the effective group ID of the calling process.

If the *path* argument refers to a regular file, the set-user-ID (S\_ISUID) and set-group-ID (S\_ISGID) bits of the file mode are cleared, if the function is successful.

If *lchown()* succeeds, the *st\_ctime* field of the file is marked for update.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EACCES Search permission is denied on a component of the path prefix.
- ELOOP Too many levels of symbolic links or prefixes.
- ENAMETOOLONG  
The length of the *path* string exceeds PATH\_MAX, or a pathname component is longer than NAME\_MAX.
- ENOENT A component of the path prefix doesn't exist, or the *path* arguments points to an empty string.
- ENOSYS The *lchown()* function isn't implemented for the filesystem specified in *path*.
- ENOTDIR A component of the path prefix isn't a directory.
- EPERM The effective user ID does not match the owner of the file, or the calling process does not have appropriate privileges.
- EROFS The named file resides on a read-only filesystem.

**Examples:**

```
/*
 * Change the ownership of a list of files
 * to the current user/group
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main(int argc, char **argv)
{
 int i;
 int ecode = 0;

 for(i = 1; i < argc; i++) {
 if(lchown(argv[i], getuid(), getgid()) == -1) {
 perror(argv[i]);
 ecode++;
 }
 }
 return(ecode);
}
```

**Classification:**

Standard Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***chmod(), chown(), errno, fchown(), fstat(), open(), stat()*

## ***lcong48()***

© 2004, QNX Software Systems Ltd.

*Initialize a sequence of pseudo-random numbers*

### **Synopsis:**

```
#include <stdlib.h>

void lcong48(unsigned short int param[7]);
```

### **Arguments:**

- param* An array of 7 short integers that are used to initialize the sequence;
- The first three entries are used to initialize the seed.
  - The next three are used to initialize the multiplicand.
  - The last entry is used to initialize the addend. You can't use values greater than **0xFFFF** as the addend.

### **Library:**

**libc**

### **Description:**

The *lcong48()* function gives you full control over the multiplicand and addend used in *drand48()*, *erand48()*, *lrand48()*, *nrand48()*, *mrnd48()*, and *jrand48()*, and the seed used in *drand48()*, *lrand48()*, and *mrnd48()*.

### **Classification:**

Standard Unix

#### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |



**See also:**

*drand48(), erand48(), jrand48(), lrand48(), mrand48(), nrand48(),  
seed48(), srand48()*

## ***ldexp()*, *ldexpf()***

© 2004, QNX Software Systems Ltd.

*Multiply a floating-point number by an integral power of 2*

### **Synopsis:**

```
#include <math.h>

double ldexp(double x,
 int exp);

float ldexp(float x,
 int exp);
```

### **Arguments:**

*x*      A floating-point number.

*exp*    The exponent of 2 to multiply *x* by.

### **Library:**

libm

### **Description:**

These functions multiply the floating-point number *x* by  $2^{exp}$ .

A range error may occur.

### **Returns:**

$x \times 2^{exp}$



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
```

```
int main(void)
{
 double value;

 value = ldexp(4.7072345, 5);
 printf("%f\n", value);

 return EXIT_SUCCESS;
}
```

produces the output:

150.631504

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*frexp(), modf()*

# ***ldiv()***

© 2004, QNX Software Systems Ltd.

*Perform division on long integers*

## **Synopsis:**

```
#include <stdlib.h>

ldiv_t ldiv(long int numer,
 long int denom);
```

## **Arguments:**

*numer*     The numerator.  
*denom*     The denominator.

## **Library:**

libc

## **Description:**

The *ldiv()* function calculates the quotient and remainder of:

*numer* ÷ *denom*

## **Returns:**

A structure of type `ldiv_t` that contains the following members:

```
typedef struct {
 long int quot; /* quotient */
 long int rem; /* remainder */
} ldiv_t;
```

## **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

void print_time(long ticks)
{
 ldiv_t sec_ticks;
 ldiv_t min_sec;
```

```
 sec_ticks = ldiv(ticks, 100);
 min_sec = ldiv(sec_ticks.quot, 60);

 printf("It took %d minutes and %d seconds.\n",
 min_sec.quot, min_sec.rem);
 }

int main(void)
{
 print_time(86712);

 return EXIT_SUCCESS;
}
```

produces the output:

```
It took 14 minutes and 27 seconds.
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*div()*

## ***lfind()***

© 2004, QNX Software Systems Ltd.

*Find an entry in a linear search table*

### **Synopsis:**

```
#include <search.h>

void * lfind(const void * key,
 const void * base,
 unsigned * num,
 unsigned width,
 int (* compare) (
 const void * element1,
 const void * element2));
```

### **Arguments:**

*key*            The object to search for.

*base*            A pointer to the first element in the table.

*num*             A pointer to an integer containing the current number of elements in the table.

*width*           The size of an element, in bytes.

*compare*        A pointer to a user-supplied function that *lfind()* calls to compare an array element with the *key*. The arguments to the comparison function are:

- *element1* — the same pointer as *key*
- *element2* — a pointer to one of the array elements.

The comparison function must return 0 if *element1* equals *element2*, or a nonzero value if the elements aren't equal.

### **Library:**

`libc`

## Description:

The *lfind()* function returns a pointer into a table indicating where an entry may be found.



---

The *lfind()* function is the same as *lsearch()*, except that if the entry isn't found, it isn't added to the table, and a NULL pointer is returned.

---

## Returns:

A pointer to the matching element, or NULL if there's no match or an error occurred.

## Examples:

This example program lets you know if the first command-line argument is a C keyword (assuming you fill in the *keywords* array with a complete list of C keywords):

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <search.h>

static const char *keywords[] = {
 "auto",
 "break",
 "case",
 "char",
 /* . */
 /* . */
 /* . */
 "while"
};

int compare(const void *, const void *);

int main(int argc, const char *argv[])
{
 unsigned num = 5;
 char *ptr;

 if(argc <= 1) return EXIT_FAILURE;

 ptr = lfind(&argv[1], keywords, &num, sizeof(char **), compare);
 if(ptr == NULL) {
 printf("'%s' is not a C keyword\n", argv[1]);

 return EXIT_FAILURE;
 } else {
```

```
 printf("'%s' is a C keyword\n", argv[1]);
 }
 return EXIT_SUCCESS;
}

/* You'll never get here. */
return EXIT_SUCCESS;
}

int compare(const void *op1, const void *op2)
{
 const char **p1 = (const char **) op1;
 const char **p2 = (const char **) op2;

 return(strcmp(*p1, *p2));
}
```

**Classification:**

Standard Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***bsearch()*, *lsearch()*



## Synopsis:

```
#include <math.h>

double lgamma(double x);

double lgamma_r(double x,
 int* signgam);

float lgammaf(float x);

float lgammaf_r(float x,
 int* signgam);
```

## Arguments:

*x*            An arbitrary number.

*signgam*    (*lgamma\_r()*, *lgammaf\_r()* only) A pointer to a location where the function can store the sign of  $\Gamma(x)$ .

## Library:

`libm`

## Description:

The *lgamma()* and *lgamma\_r()* functions return the natural log (**ln**) of the  $\Gamma$  function and are equivalent to *gamma()*. These functions return **ln** |  $\Gamma(x)$  |, where  $\Gamma(x)$  is defined as follows:

$$\int_0^{\infty} e^{-t} t^{x-1} dt$$

For  $x > 0$ :

$$\text{For } x < 1: \quad n / (\Gamma(1-x) * \sin(\pi x))$$

The results converge when  $x$  is between 0 and 1. The  $\Gamma$  function has the property:

$$\Gamma(N) = \Gamma(N-1) \times N$$

The *lgamma\** functions compute the log because the  $\Gamma$  function grows very quickly.

The *lgamma()* and *lgammaf()* functions use the external integer *signgam* to return the sign of  $\Gamma(x)$ , while *lgamma\_r()* and *lgammaf\_r()* use the user-allocated space addressed by *signgamp*.



---

The *signgam* variable isn't set until *lgamma()* or *lgammaf()* returns. For example, don't use the expression:

```
g = signgam * exp(lgamma(x));
```

to compute  $g = \Gamma(x)^x$ . Instead, compute *lgamma()* first:

```
lg = lgamma(x);
g = signgam * exp(lg);
```

---

Note that  $\Gamma(x)$  must overflow when  $x$  is large enough, underflow when  $-x$  is large enough, and generate a division by 0 exception at the singularities  $x$  a nonpositive integer.

## Returns:

$\ln |\Gamma(x)|$



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>
```

```
int main(int argc, char** argv)
{
 double a, b;

 errno = EOK;
 a = 0.5;
 b = lgamma(a);
 printf("lgamma(%f) = %f %d \n", a, b, errno);

 return(0);
}
```

produces the output:

```
lgamma(0.500000) = 0.572365 0
```

## Classification:

*lgamma()* is standard Unix; *lgamma\_r()*, *lgammaf()*, and *lgammaf\_r()* are ANSI (draft)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*gamma()*

## ***link()***

© 2004, QNX Software Systems Ltd.

*Create a link to an existing file*

### **Synopsis:**

```
#include <unistd.h>

int link(const char* existing,
 const char* new);
```

### **Arguments:**

*existing*     The path of an existing file.  
*new*         The path for the new link.

### **Library:**

`libc`

### **Description:**

The *link()* function creates a new directory entry named by *new* to refer to (that is, to be a link to) an existing file named by *existing*. The function atomically creates a new link for the existing file, and increments the link count of the file by one.



---

This implementation doesn't support using *link()* on directories or the linking of files across filesystems (different logical disks).

---

If the function fails, no link is created, and the link count of the file remains unchanged.

If *link()* succeeds, the *st\_ctime* field of the file and the *st\_ctime* and *st\_mtime* fields of the directory that contains the new entry are marked for update.

### **Returns:**

0         Success.  
-1        An error occurred (*errno* is set).

**Errors:**

|              |                                                                                                                                                                                                                                                                                 |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | A component of either path prefix denies search permission, or the link named by <i>new</i> is in a directory with a mode that denies write permission.                                                                                                                         |
| EEXIST       | The link named by <i>new</i> already exists.                                                                                                                                                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                                                                                                                  |
| EMLINK       | The number of links to the file named by <i>existing</i> would exceed LINK_MAX.                                                                                                                                                                                                 |
| ENAMETOOLONG | The length of the <i>existing</i> or <i>new</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.                                                                                                                                                       |
| ENOENT       | This error code can mean the following: <ul style="list-style-type: none"><li>• A component of either path prefix doesn't exist.</li><li>• The file named by <i>existing</i> doesn't exist.</li><li>• Either <i>existing</i> or <i>new</i> points to an empty string.</li></ul> |
| ENOSPC       | The directory that would contain the link can't be extended.                                                                                                                                                                                                                    |
| ENOSYS       | The <i>link()</i> function isn't implemented for the filesystem specified in <i>existing</i> or <i>new</i> .                                                                                                                                                                    |
| ENOTDIR      | A component of either path prefix isn't a directory.                                                                                                                                                                                                                            |
| EPERM        | The file named by <i>existing</i> is a directory.                                                                                                                                                                                                                               |
| EROFS        | The requested link requires writing in a directory on a read-only file system.                                                                                                                                                                                                  |
| EXDEV        | The link named by <i>new</i> and the file named by <i>existing</i> are on different logical disks.                                                                                                                                                                              |

**Examples:**

```
/*
 * The following program performs a rename
 * operation of argv[1] to argv[2].
 * Please note that this example, unlike the
 * library function rename(), ONLY works if
 * argv[2] doesn't already exist.
 */
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>

int main(int argc, char** argv)
{
 /* Create a link of argv[1] to argv[2].
 */
 if(link(argv[1], argv[2]) == -1) {
 perror("link");
 return(EXIT_FAILURE);
 }
 if(unlink(argv[1]) == -1) {
 perror(argv[1]);
 return(EXIT_FAILURE);
 }
 return(EXIT_SUCCESS);
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno, rename(), symlink(), unlink()*

## ***lio\_listio()***

© 2004, QNX Software Systems Ltd.

*Initiate a list of I/O requests*

### **Synopsis:**

```
#include <aio.h>

int lio_listio(int mode,
 struct aiocb* const list[],
 int nent,
 struct sigevent* sig);
```

### **Arguments:**

- mode*     The mode of operation; one of:
- LIO\_WAIT — *lio\_listio()* behaves synchronously, waiting until all I/O is completed, and ignores the *sig* argument.
  - LIO\_NOWAIT — *lio\_listio()* behaves asynchronously, returning immediately, and the signal specified by the *sig* argument is delivered to the calling process when all the I/O operations from this function complete.
- list*     An array of pointers to **aiocb** structures that specify the I/O operations that you want to initiate. The array may contain NULL pointers, which the function ignores.
- nent*     The number of entries in the *list* array. This must not exceed the system-wide limit, **\_POSIX\_AIO\_MAX**.
- sig*     NULL, or a pointer to a **sigevent** structure that specifies the signal that you want to deliver to the calling process when all of the I/O operations complete. The function ignores this argument if *mode* is LIO\_WAIT.

### **Library:**

**libc**



## Description:

The *lio\_listio()* function lets the calling process, lightweight process (LWP), or thread initiate a list of I/O requests within a single function call.

The *lio\_opcode* field of each **aio\_cb** structure in *list* specifies the operation to be performed (see `<aio.h>`):

- LIO\_READ requests *aio\_read()*.
- LIO\_WRITE requests *aio\_write()*.
- LIO\_NOP causes the list entry to be ignored.

If *mode* is LIO\_NOWAIT, *lio\_listio()* uses the **sigevent** structure pointed to by *sig* to define both the signal to be generated and how the calling process is notified when the I/O operations are complete:

- If *sig* is NULL, or the *sigevent\_signo* member of the **sigevent** structure is zero, then no signal delivery occurs. Otherwise, the signal number indicated by *sigevent\_signo* is delivered when all the requests in the list have completed.
- If *sig->sigevent\_notify* is SIGEV\_NONE, no signal is posted upon I/O completion, but the error status and the return status for the operation are set appropriately.
- If *sig->sigevent\_notify* is SIGEV\_SIGNAL, the signal specified in *sig->sigevent\_signo* is sent to the process. If the SA\_SIGINFO flag is set for that signal number, the signal is queued to the process, and the value specified in *sig->sigevent\_value* is the *si\_value* component of the generated signal.

For regular files, no data transfer occurs past the offset maximum established in the open file description associated with *aio\_cb->aio\_fildes*.

The behavior of this function is altered according to the definitions of synchronized I/O data integrity completion and synchronized I/O file integrity completion if synchronized I/O is enabled on the file

associated with *aio\_fildes*. (see the definitions of O\_DSYNC and O\_SYNC in the description of *fcntl()*.)

## Returns:

If the *mode* argument is LIO\_NOWAIT, and the I/O operations are successfully queued, *lio\_listio()* returns 0; otherwise, it returns -1, and sets *errno*.

If the *mode* argument is LIO\_WAIT, and all the indicated I/O has completed successfully, *lio\_listio()* returns 0; otherwise, it returns -1, and sets *errno*.

In either case, the return value indicates only the success or failure of the *lio\_listio()* call itself, not the status of the individual I/O requests. In some cases, one or more of the I/O requests contained in the list may fail. Failure of an individual request doesn't prevent completion of any other individual request. To determine the outcome of each I/O request, examine the error status associated with each **aio\_cb** control block. Each error status so returned is identical to that returned as a result of calling *aio\_read()* or *aio\_write()*.

## Errors:

- |        |                                                                                                                                                                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The resources necessary to queue all the I/O requests weren't available. The error status for each request is recorded in the <i>aio_error</i> member of the corresponding <b>aio_cb</b> structure, and can be retrieved using <i>aio_error()</i> .<br>The number of entries, <i>nent</i> , exceeds the system-wide limit, <code>_POSIX_AIO_MAX</code> . |
| EINVAL | The <i>mode</i> argument is invalid.<br>The value of <i>nent</i> is greater than <code>_POSIX_AIO_LISTIO_MAX</code> .                                                                                                                                                                                                                                    |
| EINTR  | A signal was delivered while waiting for all I/O requests to complete during an LIO_WAIT operation. However, the outstanding I/O requests aren't canceled. Use <i>aio_fsync()</i> to determine if any request was initiated;                                                                                                                             |

*lio\_return()* to determine if any request has completed; or *lio\_error()* to determine if any request was canceled.

- EIO One or more of the individual I/O operations failed. Use *lio\_error()* with each `aiocb` structure to determine which request(s) failed.
- ENOSYS The *lio\_listio()* function isn't supported by this implementation.

If either *lio\_listio()* succeeds in queuing all of its requests, or *errno* is set to EAGAIN, EINTR, or EIO, then some of the I/O specified from the list may have been initiated. In this event, each `aiocb` structure contains errors specific to the *read()* or *write()* function being performed:

- EAGAIN The requested I/O operation wasn't queued due to resource limitations.
- ECANCELED The requested I/O was canceled before the I/O completed due to an explicit *lio\_cancel()* request.
- EINPROGRESS The requested I/O is in progress.

The following additional error codes may be set for each `aiocb` control block:

- EOVERFLOW The *aiocbp->lio\_opcode* is LIO\_READ, the file is a regular file, *aiocbp->lio\_nbytes* is greater than 0, and the *aiocbp->lio\_offset* is before the end-of-file and is greater than or equal to the offset maximum in the open file description associated with *aiocbp->lio\_fildes*.
- EFBIG The *aiocbp->lio\_opcode* is LIO\_WRITE, the file is a regular file, *aiocbp->lio\_nbytes* is greater than 0, and the *aiocbp->lio\_offset* is greater than or equal to the offset maximum in the open file description associated with *aiocbp->lio\_fildes*.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*aio\_cancel()*, *aio\_error()*, *aio\_fsync()*, *aio\_read()*, *aio\_return()*,  
*aio\_write()*, *close()*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*,  
*execve()*, *execvp()*, *execvpe()*, *exit()*, *fcntl()*, *fork()*, *lseek()*, *read()*,  
**sigevent**, *write()*

## Synopsis:

```
#include <sys/socket.h>

int listen(int s,
 int backlog);
```

## Arguments:

- s*            The descriptor for the socket that you want to listen on. You can create a socket by calling *socket()*.
- backlog*     The maximum length that the queue of pending connections may grow to.

## Library:

**libsocket**

## Description:

The *listen()* function listens for connections on a socket and puts the socket into the LISTEN state. For connections to be accepted, you must:

- 1    Create a socket by calling *socket()*.
- 2    Indicate a willingness to accept incoming connections and a queue limit for them by calling *listen()*.
- 3    Call *accept()* to accept the connections.

If a connection request arrives with the queue full, the client may receive an error with an indication of *ECONNREFUSED*. But if the underlying protocol supports retransmission, the request may be ignored so that retries may succeed.



---

The *listen()* call applies only to SOCK\_STREAM sockets.

---

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EBADF Invalid descriptor *s*.
- EOPNOTSUPP  
The socket isn't of a type that supports the *listen()* operation.

## Classification:

Standard Unix, POSIX 1003.1-2001

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*accept()*, *connect()*, *socket()*

**Synopsis:**

```
#include <locale.h>

struct lconv * localeconv(void);
```

**Library:**

```
libc
```

**Description:**

The *localeconv()* function gets the values appropriate for formatting numeric quantities using the current locale. It returns a pointer to a **struct lconv** with the following members:

**char \* decimal\_point**

The decimal-point character used for nonmonetary quantities.

**char \* thousands\_sep**

The character used to separate groups of digits on the left of the decimal-point character formatted nonmonetary quantities.

**char \* int\_curr\_symbol**

The international currency symbol for the current locale. The first three characters contain the alphabetic international currency symbol in accordance with those specified in *ISO 4217: Codes for the Representation of Currency and Funds*. The fourth character (immediately preceding the NUL character) is the character used to separate the international currency symbol from the monetary quantity.

**char \* currency\_symbol**

The local currency symbol applicable to the current locale.

**char \* mon\_decimal\_point**

The decimal-point character used to format monetary quantities.

**char** \* *mon\_thousands\_sep*

The character used to separate groups of digits on the left of the decimal-point character in formatted monetary quantities.

**char** \* *mon\_grouping*

A string whose elements indicate the size of each group of digits in formatted monetary quantities.

**char** \* *grouping*

A string whose elements indicate the size of each group of digits in formatted nonmonetary quantities.

**char** \* *positive\_sign*

The string used to indicate a nonnegative monetary quantity.

**char** \* *negative\_sign*

The string used to indicate a negative monetary quantity.

**char** *int\_frac\_digits*

The number of fractional digits (to the right of the decimal point) to display in an internationally formatted monetary quantity.

**char** *frac\_digits*

The number of fractional digits (to the right of the decimal point) to display in a formatted monetary quantity.

**char** *p\_cs\_precedes*

Set to 1 or 0 if the *currency\_symbol* precedes or follows the value for a nonnegative monetary quantity.

**char** *p\_sep\_by\_space*

Set to 1 or 0 if the *currency\_symbol* is or isn't separated by a space from the value for a nonnegative monetary quantity.

**char** *n\_cs\_precedes*

Set to 1 or 0 if the *currency\_symbol* precedes or follows the value for a negative monetary quantity.



**char** *n\_sep\_by\_space*

Set to 1 or 0 if the *currency\_symbol* is or isn't separated by a space from the value for a negative monetary quantity.

**char** *p\_sign\_posn*

The position of the *positive\_sign* for a nonnegative monetary quantity.

**char** *n\_sign\_posn*

The position of the *positive\_sign* for a negative monetary quantity.

The *grouping* and *mon\_grouping* members have the following values:

|          |                                                                                                                                                                                   |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHAR_MAX | Perform no further grouping.                                                                                                                                                      |
| 0        | Repeat the previous element used for the remainder of the digits.                                                                                                                 |
| other    | The value is the number of digits that comprise the current group. Examine the next element to determine the size of the next group of digits (to the left of the current group). |

The *p\_sign\_posn* and *n\_sign\_posn* members have the following values:

|   |                                                                                |
|---|--------------------------------------------------------------------------------|
| 0 | Parentheses surround the quantity and <i>currency_symbol</i> .                 |
| 1 | The sign string precedes the quantity and <i>currency_symbol</i> .             |
| 2 | The sign string follows the quantity and <i>currency_symbol</i> .              |
| 3 | The sign string immediately precedes the quantity and <i>currency_symbol</i> . |
| 4 | The sign string immediately follows the quantity and <i>currency_symbol</i> .  |

**Returns:**

A pointer to the `struct lconv`.

**Examples:**

```
#include <stdio.h>
#include <locale.h>
#include <stdlib.h>

int main(void)
{
 struct lconv *lc;

 lc = localeconv();
 printf("decimal_point (%s)\n",
 lc->decimal_point);

 printf("thousands_sep (%s)\n",
 lc->thousands_sep);

 printf("int_curr_symbol (%s)\n",
 lc->int_curr_symbol);

 printf("currency_symbol (%s)\n",
 lc->currency_symbol);

 printf("mon_decimal_point (%s)\n",
 lc->mon_decimal_point);

 printf("mon_thousands_sep (%s)\n",
 lc->mon_thousands_sep);

 printf("mon_grouping (%s)\n",
 lc->mon_grouping);

 printf("grouping (%s)\n",
 lc->grouping);

 printf("positive_sign (%s)\n",
 lc->positive_sign);

 printf("negative_sign (%s)\n",
 lc->negative_sign);

 printf("int_frac_digits (%d)\n",
 lc->int_frac_digits);

 printf("frac_digits (%d)\n",
 lc->frac_digits);
}
```

```
printf("p_cs_precedes (%d)\n",
 lc->p_cs_precedes);

printf("p_sep_by_space (%d)\n",
 lc->p_sep_by_space);

printf("n_cs_precedes (%d)\n",
 lc->n_cs_precedes);

printf("n_sep_by_space (%d)\n",
 lc->n_sep_by_space);

printf("p_sign_posn (%d)\n",
 lc->p_sign_posn);

printf("n_sign_posn (%d)\n",
 lc->n_sign_posn);

return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*isalpha()*, *isascii()*, *printf()*, *scanf()*, *setlocale()*, *strcat()*, *strchr()*, *strcmp()*, *strcoll()*, *strcpy()*, *strftime()*, *strlen()*, *strpbrk()*, *strspn()*, *strtod()*, *strtok()*, *strxfrm()*

## ***localtime()***

© 2004, QNX Software Systems Ltd.

*Convert calendar time to local time*

### **Synopsis:**

```
#include <time.h>

struct tm *localtime(const time_t *timer);
```

### **Arguments:**

*timer*     A pointer to a `time_t` object that contains the calendar time that you want to convert.

### **Library:**

`libc`

### **Description:**

The *localtime()* function converts the calendar time pointed to by *timer* into local time, storing the information in a `struct tm`. Whenever you call *localtime()*, it calls *tzset()*.

You typically get a calendar time by calling *time()*. That time is Coordinated Universal Time (UTC, formerly known as Greenwich Mean Time or GMT).

The *localtime()* function places the converted time in a static structure that's reused each time you call *localtime()*. Use *localtime\_r()* if you want a thread-safe version.

You typically use the `date` command to set the computer's internal clock using Coordinated Universal Time (UTC). Use the **TZ** environment variable or `_CS_TIMEZONE` configuration string to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

**Returns:**

A pointer to the static `struct tm` containing the time information.

**Classification:**

ANSI, POSIX 1003.1

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*asctime()*, *asctime\_r()*, *clock()*, *ctime()*, *ctime\_r()*, *difftime()*, *gmtime()*, *gmtime\_r()*, *mktime()*, *localtime\_r()*, *strftime()*, *time()*, `tm`, *tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the *Neutrino User’s Guide*

## ***localtime\_r()***

© 2004, QNX Software Systems Ltd.

*Convert calendar time to local time*

### **Synopsis:**

```
#include <time.h>

struct tm* localtime_r(const time_t* timer,
 struct tm* result);
```

### **Arguments:**

*timer*     A pointer to a `time_t` object that contains the calendar time that you want to convert.

*result*    A pointer to a `tm` structure where the function can store the converted time.

### **Library:**

`libc`

### **Description:**

The *localtime\_r()* function converts the calendar time pointed to by *timer* into local time, storing the information in the `struct tm` that *result* points to. Whenever you call *localtime\_r()*, it calls *tzset()*.

You typically get a calendar time by calling *time()*. That time is Coordinated Universal Time (UTC, formerly known as Greenwich Mean Time or GMT).

You typically use the `date` command to set the computer's internal clock using Coordinated Universal Time (UTC). Use the **TZ** environment variable or `_CS_TIMEZONE` configuration string to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

**Returns:**

A pointer to *result*, the `struct tm`.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*asctime()*, *asctime\_r()*, *clock()*, *ctime()*, *ctime\_r()*, *difftime()*, *gmtime()*, *gmtime\_r()*, *localtime()*, *mktime()*, *strftime()*, *time()*, `tm`, *tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*

# ***lockf()***

© 2004, QNX Software Systems Ltd.

*Lock or unlock a section of a file*

## **Synopsis:**

```
#include <unistd.h>

int lockf(int filedes,
 int function,
 off_t size);
```

## **Arguments:**

- filedes*      The file descriptor for the file that you want to lock. Open the file with write-only permission (O\_WRONLY) or with read/write permission (O\_RDWR).
- function*     A control value that specifies the action to be taken. The permissible values (defined in <unistd.h>) are as follows:
- F\_LOCK      Lock a section for exclusive use if the section is available. A read-only lock is one of O\_RDONLY, O\_WRONLY, or O\_RDWR. An exclusive lock is one of O\_WRONLY, or O\_RDWR. (For descriptions of the locks, see *open()*).
  - F\_TEST      Test a specified section for locks obtained by other processes.
  - F\_TLOCK     Test and lock a section for exclusive use if the section is available.
  - F\_ULOCK     Remove locks from a specified section of the file.
- size*          The number of contiguous bytes that you want to lock or unlock. The section to be locked or unlocked starts at the current offset in the file and extends forward for a positive *size* or backward for a negative *size* (the preceding bytes up to but not including the current offset). If *size* is 0, the section from the current offset



through the largest possible file offset is locked (that is, from the current offset through to the present or any future end-of-file). An area need not be allocated to the file to be locked because locks may exist past the end-of-file.

## Library:

`libc`

## Description:

You can use the *lockf()* function to lock a section of a file, using advisory-mode locks. If other threads call *lockf()* to try to lock the locked file section, those calls either return an error value or block until the section becomes unlocked.

All the locks for a process are removed when the process terminates. Record locking with *lockf()* is supported for regular files and may be supported for other files.

The sections locked with `F_LOCK` or `F_TLOCK` may in whole or in part, contain or be contained by a previously locked section for the same process. When this occurs, or if adjacent locked sections occur, the sections are combined into a single locked section.

File locks are released on the first close by the locking process of any file descriptor for the file.

`F_UNLOCK` requests may release (wholly or in part) one or more locked sections controlled by the process. Locked sections are unlocked starting at the current file offset through *size* bytes or to the end of file if *size* is `(off_t)0`. When all of a locked section isn't released (that is, when the beginning or end of the area to be unlocked falls within a locked section), the remaining portions of that section are still locked by the process. Releasing the center portion of a locked section causes the remaining locked beginning and end portions to become two separate locked sections.

A potential for deadlock occurs if the threads of a process controlling a locked section are blocked by accessing another process's locked

section. If the system detects that deadlock could occur, *lockf()* fails with EDEADLK.

The interaction between *fcntl()* and *lockf()* locks is unspecified. Blocking on a section is interrupted by any signal.

If *size* is the maximum value of type `off_t` and the process has an existing lock of size 0 in this range (indicating a lock on the entire file), then an F\_ULOCK request is treated the same as an F\_LOCK request of size 0. Otherwise an F\_ULOCK request attempts to unlock only the requested section. Attempting to lock a section of a file that's associated with a buffered stream produces unspecified results.

### Returns:

- 0 Success.
- 1 An error occurred (*errno* is set). Existing locks aren't changed.

### Errors:

|                  |                                                                                                                                                                          |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES or EAGAIN | The <i>function</i> argument is F_TLOCK or F_TEST and the section is already locked by another process.                                                                  |
| EAGAIN           | The <i>function</i> argument is F_LOCK or F_TLOCK and the file is mapped with <i>mmap()</i> .                                                                            |
| EBADF            | The <i>fildev</i> argument isn't a valid open file descriptor; or <i>function</i> is F_LOCK or F_TLOCK and <i>fildev</i> isn't a valid file descriptor open for writing. |
| EDEADLK          | The <i>function</i> argument is F_LOCK and a deadlock is detected.                                                                                                       |
| EINTR            | A signal was caught during execution of the function.                                                                                                                    |
| EINVAL           | The <i>function</i> argument isn't one of F_LOCK, F_TLOCK, F_TEST or F_ULOCK; or <i>size</i> plus the current file offset is less than 0.                                |

|                      |                                                                                                                                                                          |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENOMEM               | The system can't allocate sufficient memory to store lock resources.                                                                                                     |
| EOPNOTSUPP or EINVAL | The implementation doesn't support the locking of files of the type indicated by <i>files</i> .                                                                          |
| E_OVERFLOW           | The offset of the first, or if <i>size</i> isn't 0 then the last, byte in the requested section can't be represented correctly in an object of type <code>off_t</code> . |

### Classification:

Standard Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*fcntl()*, *flock()*, *open()*

## ***log()*, *logf()***

© 2004, QNX Software Systems Ltd.

*Compute the natural logarithm of a number*

---

### **Synopsis:**

```
#include <math.h>

double log(double x);

float logf(float x);
```

### **Arguments:**

*x* The number that you want to compute the natural log of.

### **Library:**

libm

### **Description:**

The *log()* and *logf()* functions compute the natural logarithm (base *e*) of *x*:

$\log_e x$

A domain error occurs if *x* is negative. A range error occurs if *x* is zero.

### **Returns:**

The natural logarithm of *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
```

```
int main(void)
{
 printf("%f\n", log(.5));

 return EXIT_SUCCESS;
}
```

produces the output:

-0.693147

### **Classification:**

*log()* is ANSI; *logf()* is ANSI (draft)

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

*errno*, *exp()*, *log10()*, *pow()*

## ***log1p()*, *log1pf()***

© 2004, QNX Software Systems Ltd.

*Log(1+x)*

### **Synopsis:**

```
#include <math.h>

double log1p (double x);

float log1pf (float x);
```

### **Arguments:**

*x* The number that you want to add 1 to and compute the natural log of.

### **Library:**

libm

### **Description:**

The *log1p()* and *log1pf()* functions compute the value of  $\log(1+x)$ , where  $x > -1.0$ .

### **Returns:**

**If:** *log1p()* returns:

---

$x = \text{NaN}$  NaN

$x < -1.0$  -HUGE\_VAL, or NaN (*errno* is set to EDOM).

$x = -1.0$  -HUGE\_VAL (*errno* may be set to ERANGE).



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Classification:**

*log1p()* is standard Unix; *log1pf()* is ANSI (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ilogb(), log(), logb(), log10()*

## ***log10()*, *log10f()***

© 2004, QNX Software Systems Ltd.

*Compute the logarithm (base 10) of a number*

---

### **Synopsis:**

```
#include <math.h>

double log10(double x);

float log10f(float x);
```

### **Arguments:**

*x*     The number that you want to compute the log of.

### **Library:**

`libm`

### **Description:**

The *log10()* and *log10f()* functions compute the base 10 logarithm of *x*:

$\log_{10} x$

A domain error occurs if *x* is negative. A range error occurs if *x* is zero.

### **Returns:**

The base 10 logarithm of *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
```



```
int main(void)
{
 printf("%f\n", log10(.5));

 return EXIT_SUCCESS;
}
```

produces the output:

-0.301030

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno, exp(), log(), pow()*

## ***logb()*, *logbf()***

© 2004, QNX Software Systems Ltd.

*Compute the radix-independent exponent*

### **Synopsis:**

```
#include <math.h>

double logb (double x);

float logbf (float x);
```

### **Arguments:**

*x* The number that you want to compute the radix-independent exponent of.

### **Library:**

libm

### **Description:**

The *logb()* and *logbf()* functions compute the exponent part of *x*, which is the integral part of:

$$\log_r |x|$$

as a signed floating point value, for nonzero finite *x*, where *r* is the radix of the machine's floating point arithmetic.

### **Returns:**

The binary exponent of *x*, a signed integer converted to double-precision floating-point.

| <b>If <i>x</i> is:</b> | <b><i>logb()</i> returns:</b>                  |
|------------------------|------------------------------------------------|
| 0.0                    | -HUGE_VAL ( <i>errno</i> is set to EDOM)       |
| <0.0                   | -HUGE_VAL ( <i>errno</i> may be set to ERANGE) |
| ±infinity              | +infinity                                      |



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

## Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b;

 a = 0.5;
 b = logb(a);
 printf("logb(%f) = %f (%f = 2^%f) \n", a, b, a, b);

 return(0);
}
```

produces the output:

```
logb(0.500000) = -1.000000 (0.500000 = 2^-1.000000)
```

## Classification:

*logb()* is standard Unix; *logbf()* is ANSI (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ilogb()*, *log()*, *log10()*, *log1p()*

## Synopsis:

```
#include <unix.h>

int login_tty(int fd);
```

## Arguments:

*fd* A file descriptor that you want to use as the controlling terminal for the current process.

## Library:

libc

## Description:

The *login\_tty()* function prepares for a login on the tty *fd* (which may be a real tty device, or the slave of a pseudo-tty as returned by *openpty()*) by creating a new session, making *fd* the controlling terminal for the current process, setting *fd* to be the standard input, output, and error streams of the current process, and closing *fd*.

This function fails if *ioctl()* fails to set *fd* to the controlling terminal of the current process.

## Returns:

0 Success.  
-1 An error occurred; *errno* is set.

## Classification:

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*forkpty(), ioctl(), openpty()*

### **Synopsis:**

```
#include <setjmp.h>

void longjmp(jmp_buf env,
 int return_value);
```

### **Arguments:**

*env*                    The environment saved by the most recent call to *setjmp()*.

*return\_value*        The value that you want *setjmp()* to return.

### **Library:**

`libc`

### **Description:**

The *longjmp()* function restores the environment saved in *env* by the most recent call to the *setjmp()* function.



Using *longjmp()* to jump out of a signal handler can cause unpredictable behavior, unless the signal was generated by the *raise()* function.

---

### **Returns:**

After the *longjmp()* function restores the environment, program execution continues as if the corresponding call to *setjmp()* had just returned the value specified by *return\_value*. If the value of *return\_value* is 0, the value returned is 1.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <setjmp.h>
```

```
jmp_buf env;

void rtn(void)
{
 printf("about to longjmp\n");
 longjmp(env, 14);
}

int main(void)
{
 int ret_val = 293;

 if(0 == (ret_val = setjmp(env))) {
 printf("after setjmp %d\n", ret_val);
 rtn();
 printf("back from rtn %d\n", ret_val);
 } else {
 printf("back from longjmp %d\n", ret_val);
 }

 return EXIT_SUCCESS;
}
```

produces the following output:

```
after setjmp 0
about to longjmp
back from longjmp 14
```

## Classification:

ANSI, POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



## Caveats:

A strictly-conforming POSIX application can't assume that the *longjmp()* function is signal-safe on other platforms.



---

**WARNING:** *Don't use `longjmp()` or `siglongjmp()` to restore an environment saved by a call to `setjmp()` or `sigsetjmp()` in another thread. If you're lucky, your application will crash; if not, it'll look as if it works for a while, until random scribbling on the stack causes it to crash.*

---

## See also:

*setjmp(), siglongjmp(), sigsetjmp()*

## ***lrand48()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random nonnegative long integer*

### **Synopsis:**

```
#include <stdlib.h>

long lrand48(void);
```

### **Library:**

libc

### **Description:**

The *lrand48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a nonnegative **long** integer uniformly distributed over the interval  $[0, 2^{31}]$ .

Call one of *lcong48()*, *seed48()*, or *srand48()* to initialize the random-number generator before calling *drand48()*, *lrand48()*, or *mrnd48()*,

The *nrnd48()* function is a thread-safe version of *lrand48()*.

### **Returns:**

A pseudo-random **long** integer.

### **Classification:**

Standard Unix

#### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*drاند48(), erاند48(), jrand48(), lcong48(), mrand48(), nrand48(),  
seed48(), srand48()*

## ***lsearch()***

© 2004, QNX Software Systems Ltd.

*Perform a linear search in an array*

### **Synopsis:**

```
#include <search.h>

void * lsearch(const void * key,
 const void * base,
 unsigned * num,
 unsigned width,
 int (* compare) (
 const void * element1,
 const void * element2));
```

### **Arguments:**

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>key</i>     | The object to search for.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>base</i>    | A pointer to the first element in the table.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <i>num</i>     | A pointer to an integer containing the current number of elements in the table.                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>width</i>   | The size of an element, in bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <i>compare</i> | A pointer to a user-supplied function that <i>lsearch()</i> calls to compare an array element with the <i>key</i> . The arguments to the comparison function are: <ul style="list-style-type: none"><li>• <i>element1</i> — the same pointer as <i>key</i></li><li>• <i>element2</i> — a pointer to one of the array elements.</li></ul> The comparison function must return 0 if <i>element1</i> equals <i>element2</i> , or a nonzero value if the elements aren't equal. |

### **Library:**

`libc`

## Description:

The *lsearch()* function searches a linear table and returns a pointer into the table indicating where the entry was found.



---

If *key* isn't found, it's added to the end of the array and *num* is incremented.

---

## Returns:

A pointer to the element that was found or created, or NULL if an error occurred.

## Examples:

This program builds an array of pointers to the *argv* arguments by searching for them in an array of NULL pointers. Because none of the items will be found, they'll all be added to the *array*.

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <search.h>

int compare(const void *, const void *);

int main(int argc, const char **argv)
{
 int i;
 unsigned num = 0;

 char **array = (char **)calloc(argc, sizeof(char **));

 for(i = 1; i < argc; ++i) {
 lsearch(&argv[i], array, &num, sizeof(char **), compare);
 }

 for(i = 0; i < num; ++i) {
 printf("%s\n", array[i]);
 }

 return EXIT_SUCCESS;
}

int compare(const void *op1, const void *op2)
{
 const char **p1 = (const char **) op1;
 const char **p2 = (const char **) op2;

 return(strcmp(*p1, *p2));
}
```

Using the program above, this input:

```
one two one three four
```

produces the output:

```
one
two
three
four
```

## **Classification:**

Standard Unix

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*bsearch()*, *lfind()*

### **Synopsis:**

```
#include <sys/types.h>
#include <unistd.h>

off_t lseek(int filedes,
 off_t offset,
 int whence);

off64_t lseek64(int filedes,
 off64_t offset,
 int whence);
```

### **Arguments:**

- filedes*    The file descriptor of the file whose position you want to set.
- offset*    The relative offset from the file position determined by the *whence* argument.
- whence*    The position in the file. The possible values (defined in `<unistd.h>`) are:
- SEEK\_CUR    The new file position is computed relative to the current file position. The value of *offset* may be positive, negative or zero.
  - SEEK\_END    The new file position is computed relative to the end of the file.
  - SEEK\_SET    The new file position is computed relative to the start of the file. The value of *offset* must not be negative.

### **Library:**

`libc`

## Description:

These functions set the current file position for the file descriptor specified by *filedes* at the operating system level. File descriptors are returned by a successful execution of one of the *creat()*, *dup()*, *dup2()*, *fcntl()*, *open()* or *sopen()* functions.

An error occurs if the requested file position is before the start of the file.

If the requested file position is beyond the end of the file and data is written at this point, subsequent reads of data in the gap will return bytes whose value is equal to zero (`'\0'`) until data is actually written into the gap.

These functions don't extend the size of a file (see *chsize()*).

## Returns:

The current file position, with 0 indicating the start of the file, or -1 if an error occurs (*errno* is set).

## Errors:

|            |                                                                                                                    |
|------------|--------------------------------------------------------------------------------------------------------------------|
| EBADF      | The <i>filedes</i> argument isn't a valid file descriptor.                                                         |
| EINVAL     | The <i>whence</i> argument isn't a proper value, or the resulting file offset is invalid.                          |
| ENOSYS     | The <i>lseek()</i> function isn't implemented for the filesystem specified by <i>filedes</i> .                     |
| E_OVERFLOW | The resulting file offset is a value that can't be represented correctly in an object of type <code>off_t</code> . |
| ESPIPE     | The <i>filedes</i> argument is associated with a pipe or FIFO.                                                     |



## Examples:

Using the *lseek()* function, you can get the current file position (in fact, *tell()* is implemented this way). You can then use this value with another call to *lseek()* to reset the file position:

```
off_t file_posn;
int filedes;

/* get current file position */
file_posn = lseek(filedes, 0L, SEEK_CUR);

...

/* return to previous file position */
file_posn = lseek(filedes, file_posn, SEEK_SET);
```

If all records in the file are the same size, the position of the *n*th record can be calculated and read like this:

```
#include <sys/types.h>
#include <unistd.h>

int read_record(int filedes, long rec_num,
 int rec_size, char *buffer)
{
 if(lseek(filedes , rec_num * rec_size,
 SEEK_SET) == -1L) {
 return -1;
 }

 return(read(filedes , buffer, rec_size));
}
```

The *read\_record()* function in this example assumes records are numbered starting with zero, and that *rec\_size* contains the size of a record in the file, including any record-separator characters.

## Classification:

*lseek()* is POSIX 1003.1; *lseek64()* is for large-file support

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*chsize()*, *close()*, *creat()*, *dup()*, *dup2()*, *eof()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *fcntl()*, *fileno()*, *fstat()*, *isatty()*, *open()*, *read()*, *sopen()*, *stat()*, *tell()*, *umask()*, *write()*

### **Synopsis:**

```
#include <sys/stat.h>

int lstat(const char* path,
 struct stat* buf);

int lstat64(const char* path,
 struct stat64* buf);
```

### **Arguments:**

*path*     The path of the file or directory that you want information about.

*buf*      A pointer to a buffer where the function can store the information.

### **Library:**

`libc`

### **Description:**

These functions obtain information about the file or directory referenced in *path*. This information is placed in the structure located at the address indicated by *buf*.

The results of the *lstat()* function are the same as the results of *stat()* when used on a file that isn't a symbolic link. If the file is a symbolic link, *lstat()* returns information about the symbolic link, while *stat()* continues to resolve the pathname using the contents of the symbolic link, and returns information about the resulting file.

### **Returns:**

0        Success.

-1      An error occurred (*errno* is set).

## Errors:

See *stat()* for details.

## Examples:

```
/*
 * Iterate through a list of files, and report
 * for each if it is a symbolic link
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>
#include <unistd.h>

int main(int argc, char **argv)
{
 int ecode = 0;
 int n;
 struct stat sbuf;

 for(n = 1; n < argc; ++n) {
 if(lstat(argv[n], &sbuf) == -1) {
 perror(argv[n]);
 ecode++;
 } else if(S_ISLNK(sbuf.st_mode)) {
 printf("%s is a symbolic link\n", argv[n]);
 } else {
 printf("%s is not a symbolic link\n", argv[n]);
 }
 }
 return(ecode);
}
```

## Classification:

*lstat()* is POSIX 1003.1a; *lstat64()* is for large-file support

### **Safety**

---

Cancellation point    No

Interrupt handler    No

*continued...*

**Safety**

---

Signal handler      Yes

Thread              No

**See also:**

*errno, fstat(), readlink(), stat()*

## ***ltoa()*, *lltoa()***

© 2004, QNX Software Systems Ltd.

*Convert a long integer into a string, using a given base*

### **Synopsis:**

```
#include <stdlib.h>

char* ltoa(long value,
 char* buffer,
 int radix);

char* lltoa(int64_t value,
 char* buffer,
 int radix);
```

### **Arguments:**

*value*     The value to convert into a string.

*buffer*    A buffer in which the function stores the string. The size of the buffer must be at least 33 bytes when converting values in base 2 (binary).

*radix*     The base to use when converting the number. This value must be in the range:

$$2 \leq \textit{radix} \leq 36$$

If the value of *radix* is 10, and *value* is negative, then a minus sign is prepended to the result.

### **Library:**

`libc`

### **Description:**

The *ltoa()* and *lltoa()* functions convert the given long integer *value* into the equivalent string in base *radix* notation, storing the result in the character array pointed to by *buffer*. A NUL character is appended to the result.

**Returns:**

A pointer to the result.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

void print_value(long value)
{
 int base;
 char buffer[33];

 for(base = 2; base <= 16; base = base + 2) {
 printf("%2d %s\n", base,
 ltoa(value, buffer, base));
 }
}

int main(void)
{
 print_value(12765);

 return EXIT_SUCCESS;
}
```

produces the output:

```
 2 11000111011101
 4 3013131
 6 135033
 8 30735
10 12765
12 7479
14 491b
16 31dd
```

**Classification:**

*ltoa()* is QNX 4; *lltoa()* is Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atoi()*, *atol()*, *ltoa()*, *sscanf()*, *strtol()*, *strtoul()*, *ultoa()*, *utoa()*



**Synopsis:**

```
#include <sys/types.h>
#include <unistd.h>

off_t ltrunc(int fildes,
 off_t offset,
 int whence);
```

**Arguments:**

- |               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>fildes</i> | The file descriptor of the file that you want to truncate.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <i>offset</i> | The relative offset from the file position determined by the <i>whence</i> argument.                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>whence</i> | The position in the file. The possible values (defined in <code>&lt;unistd.h&gt;</code> ) are:<br><br>SEEK_CUR The new file position is computed relative to the current file position. The value of <i>offset</i> may be positive, negative or zero.<br>SEEK_END The new file position is computed relative to the end of the file.<br>SEEK_SET The new file position is computed relative to the start of the file. The value of <i>offset</i> must not be negative. |

**Library:**

`libc`

**Description:**

The *ltrunc()* function attempts to truncate the file at a specified position. The file, referenced by the open file descriptor *fildes*, must have been opened `O_WRONLY` or `O_RDWR`. The truncation point is calculated using the value of *offset* as a relative offset from a file position determined by the value of the argument *whence*. The value

of *offset* may be negative, although a negative truncation point (one before the beginning of the file) is an error.



---

The *ltrunc()* function ignores advisory locks that may have been set by *fcntl()*.

---

The calculated truncation point, if within the existing bounds of the file, determines the new file size; all data after the truncation point no longer exists. If the truncation point is past the existing end of file, the file size isn't changed. An error occurs if you attempt to truncate before the beginning of the file (that is, a negative truncation point).



---

The current seek position isn't changed by this function under any circumstance, including the case where the current seek position is beyond the truncation point.

---

## Returns:

Upon successful completion, this function returns the new file size. If a truncation point beyond the existing end of file was specified, the existing file size is returned, and the file size remains unchanged. Otherwise, *ltrunc()* returns a value of -1 and sets *errno* to indicate the error. The file size remains unchanged in the event of an error.

## Errors:

|        |                                                                                                                                            |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF  | The <i>fdes</i> argument isn't a valid file descriptor, open for writing.                                                                  |
| EINVAL | The <i>whence</i> argument isn't a proper value, or the resulting file size would be invalid.                                              |
| ENOSYS | An attempt was made to truncate a file of a type that doesn't support truncation (for example, a file associated with the device manager). |
| ESPIPE | The <i>fdes</i> argument is associated with a pipe or FIFO.                                                                                |

**Examples:**

```

#include <stdio.h>
#include <fcntl.h>
#include <sys/types.h>
#include <unistd.h>

char buffer[1000];

int main(void)
{
 int fd, stat;

 fd = open("test", O_CREAT | O_RDWR, 0666);
 if(fd == -1) {
 fprintf(stderr, "Open error\n");
 exit(-1);
 }

 /* Create a 1000-byte file */
 write(fd, buffer, 1000);

 /* Seek back to offset 500 and truncate the file */
 if(ltrunc(fd, 500, SEEK_SET) == -1) {
 fprintf(stderr, "ltrunc error\n");
 exit(-1);
 }
 close(fd);
 fd = open("test", O_CREAT | O_RDWR, 0666);
 printf("File size = %ld\n",
 lseek(fd, 0, SEEK_END));
 close(fd);

 return 0;
}

```

**Classification:**

QNX 4

**Safety**

Cancellation point Yes

Interrupt handler No

*continued...*

## **Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

## **Caveats:**

The *ltrunc()* function *isn't portable*, and shouldn't be used in new code. Use *ftruncate()* instead.

## **See also:**

*errno*, *ftruncate()*, *lseek()*

## Synopsis:

```
int main(void);

int main(int argc,
 const char *argv[]);

int main(int argc,
 const char *argv[],
 char *envp[]);
```

## Arguments:

The arguments depend on which form of *main()* that you use.

- argc*    The number of entries in the *argv* array.
- argv*    An array of pointers to strings that contain the arguments to the program.
- envp*    An array of pointers to strings that define the environment for the program.

## Library:

`libc`

## Description:

The *main()* function is supplied by the user and is where program execution begins. The command line to the program is broken into a sequence of tokens separated by blanks, and are passed to *main()* as an array of pointers to character strings in *argv*. The number of arguments found is passed in the parameter *argc*.

The *argv[0]* argument is a pointer to a character string containing the program name. The last element of the array pointed to by *argv* is NULL (*argv[argc]* is NULL). Arguments containing blanks can be passed to *main()* by enclosing them in quote characters (which are

removed from that element in the *argv* vector). See your shell's documentation for details.

The *envp* argument points to an array of pointers to character strings that are the environment strings for the current process. This value is identical to the *environ* variable, which is defined in the `<stdlib.h>` header file.

## Returns:

A value back to the calling program (usually the operating system).

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv)
{
 int i;
 for(i = 0; i < argc; ++i) {
 printf("argv[%d] = %s\n", i, argv[i]);
 }

 return EXIT_SUCCESS;
}
```

produces the output:

```
argv[0] = ./myppgm
argv[1] = hhhhh
argv[2] = another arg
```

when the program `myppgm` is run from the shell:

```
$./myppgm hhhhh "another arg"
```

## Classification:

ANSI

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

*abort()*, *atexit()*, *\_argc*, *\_argv*, *\_auxv*, *close()*, *execl()*, *execle()*,  
*execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *\_exit()*,  
*exit()*, *getenv()*, *putenv()*, *sigaction()*, *signal()*, *spawn()*, *spawnl()*,  
*spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*,  
*spawnvp()*, *spawnvpe()*, *system()*, *wait()*, *waitpid()*

## ***mallinfo()***

© 2004, QNX Software Systems Ltd.

*Get memory allocation information*

---

### **Synopsis:**

```
#include <malloc.h>

struct mallinfo mallinfo (void);
```

### **Library:**

libc

### **Description:**

The *mallinfo()* function returns memory-allocation information in the form of a **struct mallinfo**:

```
struct mallinfo {
 int arena; /* size of the arena */
 int ordblks; /* number of big blocks in use */
 int smlblks; /* number of small blocks in use */
 int hblks; /* number of header blocks in use */
 int hblkhd; /* space in header block headers */
 int usmlbks; /* space in small blocks in use */
 int fsmblks; /* memory in free small blocks */
 int uordblks; /* space in big blocks in use */
 int fordblks; /* memory in free big blocks */
 int keepcost; /* penalty if M_KEEP is used
 -- not used */
};
```

### **Returns:**

A **struct mallinfo**.

### **Classification:**

ANSI

#### **Safety**

---

Cancellation point No

*continued...*



**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*calloc(), free(), malloc(), realloc()*

# ***malloc()***

*Allocate memory*

© 2004, QNX Software Systems Ltd.

## **Synopsis:**

```
#include <stdlib.h>

void* malloc(size_t size);
```

## **Arguments:**

*size*     The number of bytes to allocate.

## **Library:**

**libc**

## **Description:**

The *malloc()* function allocates a buffer of *size* bytes.



---

This function allocates memory in blocks of *\_amblksiz* bytes (a global variable defined in **<stdlib.h>**).

---

## **Returns:**

A pointer to the start of the allocated memory, or NULL if an error occurred (*errno* is set).

## **Errors:**

ENOMEM     Not enough memory.  
EOK         No error.

## **Examples:**

```
#include <stdlib.h>

int main(void)
{
 char* buffer;

 buffer = (char*)malloc(80);
```

```
 if(buffer != NULL) {
 /* do something with the buffer */
 ...

 free(buffer);
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

In QNX 4, nothing is allocated when you *malloc()* 0 bytes. Be careful if your code is ported between QNX 4 and QNX Neutrino.

## See also:

*calloc()*, *free()*, *realloc()*, *sbrk()*

# ***mallocpt()***

© 2004, QNX Software Systems Ltd.

*Control the memory allocation*

## **Synopsis:**

```
#include <malloc.h>

int mallocpt(int cmd,
 int value);
```

## **Arguments:**

*cmd* Supported arguments are:

- MALLOC\_ARENA\_SIZE — Sets the size of the requests to the system for additional core memory. The *value* argument is rounded up to the nearest page size. An argument of 0 returns the current arena size — any other argument sets the arena size, and returns the previous value.
- MALLOC\_MONOTONIC\_GROWTH — Changes the strategy for growing blocks using the *realloc()* function. Setting *value* to a nonzero number causes resize requests that don't fit in the existing block to grow the block by a minimum of 100%.

These arguments are silently ignored. They're included for compatibility reasons:

- M\_GRAIN
- M\_KEEP
- M\_MMAP\_THRESHOLD
- M\_MMAP\_MAX
- M\_MXFAST
- M\_NLBLKS
- M\_TOP\_PAD
- M\_TRIM\_THRESHOLD

*value* The allocation size.

**Library:**

libc

**Description:**

The *mallopt()* function controls the memory allocation.

**Returns:**

Unless otherwise specified, the *mallopt()* returns 0 on success, or -1 if an error occurs (*errno* is set).

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*calloc()*, *free()*, *mallinfo()*, *malloc()*, *realloc()*

## ***max()***

© 2004, QNX Software Systems Ltd.

*Return the greater of two numbers*

---

### **Synopsis:**

```
#include <stdlib.h>

#define max(a,b) ...
```

### **Arguments:**

*a,b*    The numbers that you want to get the greater of.

### **Library:**

`libc`

### **Description:**

The *max()* function returns the greater of two values.



---

The *max()* function is for C programs only. For C+ and C++ programs, use the *\_max()* or *\_min()* macros.

---

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int a;

 a = max(1, 10);
 printf("The value is: %d\n", a);
 return EXIT_SUCCESS;
}
```

### **Classification:**

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*min()*

## ***mblen()***

© 2004, QNX Software Systems Ltd.

*Count the bytes in a multibyte character*

### **Synopsis:**

```
#include <stdlib.h>

int mblen(const char * s,
 size_t n);
```

### **Arguments:**

*s* NULL (see below), or a pointer to a multibyte character.  
*n* The maximum number of bytes that you want to count.

### **Library:**

`libc`

### **Description:**

The *mblen()* function counts the number of bytes in the multibyte character pointed to by *s*, to a maximum of *n* bytes.

The *mbrlen()* function is a restartable version of *mblen()*.

### **Returns:**

- If *s* is NULL, *mblen()* determines whether or not the character encoding is state-dependent:
  - 0 The *mblen()* function uses locale-specific multibyte character encoding that's not state-dependent.
  - ≠ 0 Character is state-dependent.
- If *s* isn't NULL:
  - 0 *s* points to the null character.
  - 1 The next *n* bytes don't form a valid multibyte character.
  - > 0 The number of bytes that comprise the multibyte character (if the next *n* or fewer bytes form a valid multibyte character).



**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int len;
 char *mbs = "string";

 printf("Character encodings do ");
 if(!mblen(NULL, 0)) {
 printf("not ");
 }
 printf("have state-dependent \nencoding.\n");

 len = mblen("string", 6);
 if(len != -1) {
 mbs[len] = '\0';
 printf("Multibyte char '%s' (%d)\n", mbs, len);
 }

 return EXIT_SUCCESS;
}
```

This produces the output:

```
Character encodings do not have state-dependent
encoding.
Multibyte char 's' (1)
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“Character manipulation functions” and “Wide-character functions”  
in the summary of functions chapter.

**Synopsis:**

```
#include <wchar.h>

size_t mbrlen(const char * s,
 size_t n,
 mbstate_t * ps);
```

**Arguments:**

- s*      A pointer to a multibyte character.
- n*      The maximum number of bytes that you want to count.
- ps*     An internal pointer that lets *mbrlen()* be a restartable version of *mblen()*; if *ps* is NULL, *mbrlen()* uses its own internal variable. You can call *mbsinit()* to determine the status of this variable.

**Library:**

`libc`

**Description:**

The *mbrlen()* function counts the bytes in the multibyte character pointed to by *s*, to a maximum of *n* bytes.

**Returns:**

- `(size_t)-2`    The resulting conversion state indicates an incomplete multibyte character after all *n* characters were converted.
- `(size_t)-1`    The function detected an encoding error before completing the next multibyte character, in which case the function *errno* to EILSEQ and leaves the resulting conversion state undefined.
- 0             The next completed character is a null character, in which case the resulting conversion state is the initial conversion state.

*x*                    The number of bytes needed to complete the next multibyte character, in which case the resulting conversion state indicates that *x* bytes have been converted.

**Errors:**

EILSEQ            Invalid character sequence.  
EINVAL            The *ps* argument points to an invalid object.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*  
“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## Synopsis:

```
#include <wchar.h>

size_t mbrtowc(wchar_t * pwc,
 const char * s,
 size_t n,
 mbstate_t * ps);
```

## Arguments:

- pwc* A pointer to a `wchar_t` object where the function can store the wide character.
- s* A pointer to the multibyte character that you want to convert.
- n* The maximum number of bytes in the multibyte character to convert.
- ps* An internal pointer that lets *mbrtowc()* be a restartable version of *mbtowc()*; if *ps* is NULL, *mbrtowc()* uses its own internal variable.
- You can call *mbsinit()* to determine the status of this variable.

## Library:

`libc`

## Description:

The *mbrtowc()* function converts single multibyte characters pointed to by *s* into wide characters pointed to by *pwc*, to a maximum of *n* bytes (not characters).

This function is affected by LC\_TYPE.

**Returns:**

- (`size_t`)-2 After converting all *n* characters, the resulting conversion state indicates an incomplete multibyte character.
- (`size_t`)-1 The function detected an encoding error before completing the next multibyte character; the function sets *errno* to EILSEQ and leaves the resulting conversion state undefined.
- 0 The next completed character is a null character; the resulting conversion state is the same as the initial one.
- x* The number of bytes needed to complete the next multibyte character, in which case the resulting conversion state indicates that *x* bytes have been converted.

**Errors:**

- EILSEQ Invalid character sequence.
- EINVAL The *ps* argument points to an invalid object.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*

“Multibyte character functions,” “Stream I/O functions,” and  
“Wide-character functions” in the summary of functions chapter.

## ***mbsinit()***

© 2004, QNX Software Systems Ltd.

*Determine the status of the conversion object used for restartable mb\*() functions*

### **Synopsis:**

```
#include <wchar.h>

int mbsinit(const mbstate_t * ps);
```

### **Arguments:**

*ps* A pointer to the conversion object that you want to test.

### **Library:**

`libc`

### **Description:**

The following functions use an object of type `mbstate_t` so that they can be restarted:

- *mbrlen()*
- *mbrtowc()*
- *mbsrtowcs()*
- *mbstowcs()*
- *wcsrtombs()*
- *wcrtomb()*

The *mbsinit()* function determines whether or not the `mbstate_t` object pointed to by *ps* describes an initial conversion state.



If the object doesn't describe an initial conversion state, it isn't safe for you to use it in one of the above functions, other than the one you've already used it in.

---



## Returns:

A nonzero value if *ps* is NULL or *\*ps* describes an initial conversion state; otherwise zero.

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*mbrlen()*, *mbrtowc()*, *mbsrtowcs()*, *mbstowcs()*, *wcsrtombs()*, *wcrtomb()*

“Multibyte character functions” and “Wide-character functions” in the summary of functions chapter.

## ***mbsrtowcs()***

© 2004, QNX Software Systems Ltd.

*Convert a multibyte-character string into a wide-character string (restartable)*

### **Synopsis:**

```
#include <wchar.h>

size_t mbsrtowcs(wchar_t * dst,
 const char ** src,
 size_t n,
 mbstate_t * ps);
```

### **Arguments:**

*dst* A pointer to a buffer where the function can store the wide-character string.

*src* The string of multibyte characters that you want to convert.

*n* The maximum number of bytes that you want to convert.

*ps* An internal pointer that lets *mbsrtowcs()* be a restartable version of *mbstowcs()*; if *ps* is NULL, *mbsrtowcs()* uses its own internal variable.

You can call *mbsinit()* to determine the status of this variable.

### **Library:**

`libc`

### **Description:**

The *mbsrtowcs()* function converts a string of multibyte characters pointed to by *src* into the corresponding wide characters pointed to by *dst*, to a maximum of *n* bytes, including the terminating NULL character.

The function converts each character as if by a call to *mbtowc()* and stops early if:

- A sequence of bytes doesn't conform to a valid character

Or:

- Converting the next character would exceed the limit of  $n$  total bytes.

This function is affected by LC\_TYPE.

## Returns:

(`size_t`)-1      Failure; invalid wide-character code.  
 $x$                 Success; the number of total bytes successfully converted, not including the terminating NULL byte.

## Errors:

EILSEQ      Invalid character sequence.  
EINVAL      The  $ps$  argument points to an invalid object.

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*

“Multibyte character functions” and “Wide-character functions” in the summary of functions chapter.

## ***mbstowcs()***

© 2004, QNX Software Systems Ltd.

*Convert a multibyte-character string into a wide-character string*

### **Synopsis:**

```
#include <stdlib.h>

size_t mbstowcs(wchar_t * pwcs,
 const char * s,
 size_t n);
```

### **Arguments:**

- pwcs*     A pointer to a buffer where the function can store the wide-character string.
- s*        The string of multibyte characters that you want to convert.
- n*        The maximum number of bytes that you want to convert.

### **Library:**

`libc`

### **Description:**

The *mbstowcs()* function converts a sequence of multibyte characters pointed to by *s* into their corresponding wide-character codes pointed to by *pwcs*, to a maximum of *n* bytes. It doesn't convert any multibyte characters beyond a NULL character.

This function is affected by LC\_TYPE.

The *mbsrtowcs()* function is a restartable version of *mbstowcs()*.

### **Returns:**

The number of array elements modified, not including the terminating zero code, if present, or `(size_t) -1` if an invalid multibyte character was encountered.

**Errors:**

EILSEQ     Invalid character sequence.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *wc = "string";
 wchar_t wbuffer[50];
 int i, len;

 len = mbstowcs(wbuffer, wc, 50);
 if(len != -1) {
 wbuffer[len] = '\0';
 printf("%s(%d)\n", wc, len);

 for(i = 0; i < len; i++) {
 printf("%4.4x", wbuffer[i]);
 }

 printf("\n");
 }

 return EXIT_SUCCESS;
}
```

This produces the output:

```
string(6)
/0073/0074/0072/0069/006e/0067
```

**Classification:**

ANSI

**Safety**

Cancellation point    No

Interrupt handler      No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*errno*

“Multibyte character functions” and “Wide-character functions” in the summary of functions chapter.

## Synopsis:

```
#include <stdlib.h>

int mbtowc(wchar_t * pwc,
 const char * s,
 size_t n);
```

## Arguments:

- pwc* A pointer to a `wchar_t` object where the function can store the wide character.
- s* NULL (see below), or a pointer to the multibyte character that you want to convert.
- n* The maximum number of bytes in the multibyte character to convert.

## Library:

`libc`

## Description:

The `mbtowc()` function converts a single multibyte character pointed to by *s* into a wide-character code pointed to by *pwc*, to a maximum of *n* bytes. The function stops early if it encounters the NULL character.

This function is affected by `LC_TYPE`.

The `mbrtowc()` function is a restartable version of `mbtowc()`.

## Returns:

- If *s* is NULL:
  - 0 The `mbtowc()` function uses UTF-8 multibyte character encoding that's not state-dependent.
  - ≠ 0 Everything else.

- If *s* isn't NULL:
  - 0 The *s* argument points to the NUL character.
  - > 0 The number of bytes that comprise the multibyte character, to a maximum of MB\_CUR\_MAX (if the next *n* or fewer bytes form a valid multibyte character).
  - 1 The next *n* bytes don't form a valid multibyte character; *errno* is set.

## Errors:

EILSEQ Invalid character sequence.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *wc = "string";
 wchar_t wbuffer[10];
 int i, len;

 printf("State-dependent encoding? ");
 if(mbtowc(wbuffer, NULL, 0)) {
 printf("Yes\n");
 } else {
 printf("No\n");
 }

 len = mbtowc(wbuffer, wc, 2);
 wbuffer[len] = '\0';
 printf("%s(%d)\n", wc, len);

 for(i = 0; i < len; i++) {
 printf("%4.4x", wbuffer[i]);
 }

 printf("\n");

 return EXIT_SUCCESS;
}
```

This produces the output:



State-dependent encoding? No  
string(1)  
/0073

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*

“Multibyte character functions” and “Wide-character functions” in the summary of functions chapter

# ***mcheck()***

© 2004, QNX Software Systems Ltd.

*Enable memory allocation routine consistency checks*

## **Synopsis:**

```
#include <malloc.h>

int mcheck(
 void (* abort_fn) (enum mcheck_status status));
```

## **Arguments:**

*abort\_fn* A pointer to the callback function to invoke when an inconsistency in the memory-allocation routines is found, or NULL if you want to use the default callback routine.

The argument to the callback routine is one of the values of the **mcheck\_status** enumeration described in the documentation for *mprobe()*.

The default abort callback prints a message to *stderr* and aborts the application.

## **Library:**

**libc**

## **Description:**

The *mcheck()* function enables consistency checks within the memory allocation routines. When enabled, consistency checks are periodically performed on allocated memory blocks as blocks are allocated or freed. If an inconsistency is found, the *abort* callback is called with the status identifying the type of inconsistency found.



---

Consistency checking isn't performed on blocks that you allocated before calling *mcheck()*.

---

The level of checking provided depends on which version of the allocator you've linked the application with:

- C library — minimal consistency checking.

- Nondebug version of the `malloc` library — a slightly greater level of consistency checking.
- Debug version of the `malloc` library — extensive consistency checking, with tuning available through the use of the `mallopt()` function.

### Returns:

- 1     Checking is already enabled.
- 0     Checking wasn't already enabled.

### Classification:

QNX Neutrino

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*mallopt()*, *mprobe()*

Heap Analysis in the *Programmer's Guide*

## ***mem\_offset(), mem\_offset64()***

© 2004, QNX Software Systems Ltd.

*Get the offset of a mapped typed memory block*

### **Synopsis:**

```
#include <sys/mman.h>

int mem_offset(const void * addr,
 int fd,
 size_t length,
 off_t * offset,
 size_t * contig_len);

int mem_offset64(const void * addr,
 int fd,
 size_t length,
 off64_t * offset,
 size_t * contig_len);
```

### **Arguments:**

|                   |                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>addr</i>       | The address of the memory block whose offset and contiguous length you want to get.                                                                                                                                                    |
| <i>fd</i>         | The file descriptor that identifies the typed memory object. This must be the descriptor that you used (in a call to <i>mmap()</i> ) to establish the mapping that contains <i>addr</i> .                                              |
| <i>length</i>     | The length of the block of memory that you want the offset for.                                                                                                                                                                        |
| <i>offset</i>     | A pointer to a location where the function can store the offset of the memory block.                                                                                                                                                   |
| <i>contig_len</i> | A pointer to a location where the function can store either <i>length</i> or the length of the largest contiguous block of typed memory that's currently mapped to the calling process starting at <i>addr</i> , whichever is smaller. |

**Library:**`libc`**Description:**

The *mem\_offset()* and *mem\_offset64()* functions set the variable pointed to by *offset* to the offset (or location), within a typed memory object, of the memory block currently mapped at *addr*.

If you use the *offset* and *contig\_len* values obtained from calling *mem\_offset()* in a call to *mmap()* with a file descriptor that refers to the same memory pool as *fd* (either through the same port or through a different port), the memory region that's mapped must be exactly the same region that was mapped at *addr* in the address space of the process that called *mem\_offset()*.

**QNX extension**

If you specify *fd* as NOFD, *offset* is the offset into */dev/mem* of *addr* (i.e. its physical address). If the memory object specified by *fd* isn't a typed memory object, or specified as NOFD, the call fails.



---

If the physical address is not a valid *offset* value, *mem\_offset()* will fail with *errno* set to E2BIG. This is typically the case with many ARM systems, and you should use *mem\_offset64()* to get the physical address.

---

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EACCES The process hasn't mapped memory at the given address *addr*.
- EBADF Invalid open file descriptor *fd*.

- EINVAL      The file descriptor *fdes* doesn't correspond to the memory object mapped at *addr*.
- ENODEV     The file descriptor *fdes* isn't connected to a memory object supported by this function.
- ENOSYS     The *mem\_offset()* function isn't supported by this implementation.

## Examples:

```
#include <unistd.h>
#include <errno.h>
#include <sys/types.h>
#include <sys/mman.h>

paddr_t mphys(void *addr) {
 off64_t offset;

 if(mem_offset64(addr, NOFD, 1, &offset, 0) == -1) {
 return -1;
 }
 return offset;
}
```

## Classification:

*mem\_offset()* is POSIX 1003.1j (draft); *mem\_offset64()* is for large-file support

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*mmap(), posix\_mem\_offset(), posix\_mem\_offset64()*

# ***memalign()***

© 2004, QNX Software Systems Ltd.

*Allocate aligned memory*

---

## **Synopsis:**

```
#include <malloc.h>

void *memalign(size_t alignment,
 size_t size);
```

## **Arguments:**

*alignment*      The alignment that you want to use for the memory. This must be a multiple of **size( void \*)**.

*size*            The amount of memory you want to allocate, in bytes.

## **Library:**

**libc**

## **Description:**

The *memalign()* function allocates *size* bytes aligned on a boundary specified by *alignment*.

## **Returns:**

A pointer to the allocated block, or NULL if an error occurred (*errno* is set).

## **Errors:**

EINVAL          The value of *alignment* isn't a multiple of **size( void \*)**.

ENOMEM         There's insufficient memory available with the requested alignment.



**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*free()*, *malloc()*, *posix\_memalign()*

## ***memcpy()***

© 2004, QNX Software Systems Ltd.

*Copy bytes between buffers until a given byte is found*

### **Synopsis:**

```
#include <string.h>

void* memcpy(void* dest,
 const void* src,
 int c,
 size_t cnt);
```

### **Arguments:**

*dest*    A pointer to where you want the function to copy the data.

*src*     A pointer to the buffer that you want to copy data from.

*c*       The value that you want to stop copying at.

*cnt*     The maximum number of bytes to copy.

### **Library:**

`libc`

### **Description:**

The *memcpy()* function copies bytes from *src* to *dest*, up to and including the first occurrence of the character *c*, or until *cnt* bytes have been copied, whichever comes first.

### **Returns:**

A pointer to the byte in *dest* following the character *c*, if one is found and copied; otherwise, NULL.

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

char* msg = "This is the string: not copied";
```

```
int main(void)
{
 char buffer[80];

 memset(buffer, '\0', 80);
 memcpy(buffer, msg, ':', 80);

 printf("%s\n", buffer);

 return EXIT_SUCCESS;
}
```

produces the output:

This is the string:

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memchr()*, *memcmp()*, *memcpy()*, *memicmp()*, *memmove()*, *memset()*

## ***memchr()***

© 2004, QNX Software Systems Ltd.

*Find the first occurrence of a character in a buffer*

### **Synopsis:**

```
#include <string.h>

void* memchr(void* buf,
 int ch,
 size_t length);
```

### **Arguments:**

*buf*            The buffer that you want to search.

*ch*             The character that you're looking for.

*length*        The number of bytes to search in the buffer.

### **Library:**

libc

### **Description:**

The *memchr()* function locates the first occurrence of *ch* (converted to an **unsigned char**) in the first *length* bytes of the buffer pointed to by *buf*.

### **Returns:**

A pointer to the located character, or NULL if *ch* couldn't be found.

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];
 char* where;

 strcpy(buffer, "video x-rays");
```

```
where = (char *) memchr(buffer, 'x', 6);
if(where == NULL) {
 printf("'x' not found\n");
} else {
 printf("%s\n", where);
}

where = (char *) memchr(buffer, 'r', 9);
if(where == NULL) {
 printf("'r' not found\n");
} else {
 printf("%s\n", where);
}

return EXIT_SUCCESS;
}
```

produces the output:

```
'x' not found
rays
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memccpy()*, *memcmp()*, *memcpy()*, *memicmp()*, *memmove()*, *memset()*  
*strchr()*, *strrchr()*, *wmemchr()*, *wmemcmp()*, *wmemcpy()*,  
*wmemmove()*, *wmemset()*

## ***memcmp()***

© 2004, QNX Software Systems Ltd.

*Compare the bytes in two buffers*

---

### **Synopsis:**

```
#include <string.h>

int memcmp(const void* s1,
 const void* s2,
 size_t length);
```

### **Arguments:**

*s1, s2*     Pointers to the buffers that you want to compare.  
*length*     The number of bytes that you want to compare.

### **Library:**

libc

### **Description:**

The *memcmp()* function compares *length* bytes of the buffer pointed to by *s1* to the buffer pointed to by *s2*.

### **Returns:**

< 0     The object pointed to by *s1* is less than the object pointed to by *s2*.  
0       The object pointed to by *s1* is equal to the object pointed to by *s2*.  
> 0     The object pointed to by *s1* is greater than the object pointed to by *s2*.

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
```

```
char buffer[80];
int retval;

strcpy(buffer, "World");

retval = memcmp(buffer, "hello", 5);
if(retval < 0) {
 printf("Less than\n");
} else if(retval == 0) {
 printf("Equal to\n");
} else {
 printf("Greater than\n");
}

return EXIT_SUCCESS;
}
```

produces the output:

```
Less than
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memccpy()*, *memchr()*, *memcpy()*, *memicmp()*, *memmove()*, *memset()*

# ***memcpy()***

© 2004, QNX Software Systems Ltd.

*Copy bytes from one buffer to another*

## **Synopsis:**

```
#include <string.h>

void* memcpy(void* dst,
 const void* src,
 size_t length);
```

## **Arguments:**

*dst*        A pointer to where you want the function to copy the data.

*src*        A pointer to the buffer that you want to copy data from.

*length*     The number of bytes to copy.

## **Library:**

libc

## **Description:**

The *memcpy()* function copies *length* bytes from the buffer pointed to by *src* into the buffer pointed to by *dst*.



---

Copying overlapping buffers isn't guaranteed to work; use *memmove()* to copy buffers that overlap.

---

## **Returns:**

A pointer to the destination buffer (that is, the value of *dst*).

## **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];
```



```
memcpy(buffer, "Hello", 5);
buffer[5] = '\0';
printf("%s\n", buffer);

return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memccpy()*, *memchr()*, *memcmp()*, *memicmp()*, *memmove()*, *memset()*

## ***memcpyv()***

© 2004, QNX Software Systems Ltd.

*Copy a given number of structures*

### **Synopsis:**

```
#include <string.h>

size_t memcpyv(const struct iovec *dst,
 int dparts,
 int doff,
 const struct iovec *src,
 int sparts,
 int soff);
```

### **Arguments:**

|               |                                                                          |
|---------------|--------------------------------------------------------------------------|
| <i>dst</i>    | An array of <b>iovec</b> structures that you want to copy the data to.   |
| <i>dparts</i> | The number of entries in the <i>dst</i> array.                           |
| <i>doff</i>   | The offset into the <i>dst</i> array at which to start copying.          |
| <i>src</i>    | An array of <b>iovec</b> structures that you want to copy the data from. |
| <i>sparts</i> | The number of entries in the <i>src</i> array.                           |
| <i>soff</i>   | The offset into the <i>src</i> array at which to start copying.          |

### **Library:**

**libc**

### **Description:**

The function *memcpyv()* copies data pointed to by the *src* I/O vector, starting at offset *soff*, to *dst* structures, starting at offset *doff*. The number of I/O vector parts copied is specified in *sparts* and *dparts*.

**Returns:**

The number of bytes copied.

**Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 const struct iovec *dest, *source;
 int dparts, doffset, sparts, soffset;
 size_t nbytes;

 nbytes = memcpyv (dest, dparts, doffset,
 source, sparts, soffset);
 printf ("The number of bytes copied is %d. \n", nbytes);

 return EXIT_SUCCESS;
}
```

**Classification:**

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memccpy()*, *memcpy()*

## ***memicmp()***

© 2004, QNX Software Systems Ltd.

*Compare two buffers, ignoring case*

### **Synopsis:**

```
#include <string.h>

int memicmp(const void* s1,
 const void* s2,
 size_t length);
```

### **Arguments:**

*s1, s2*     Pointers to the buffers that you want to compare.

*length*     The number of bytes that you want to compare.

### **Library:**

libc

### **Description:**

The *memicmp()* function compares (case insensitive) *length* bytes of the buffer pointed to by *s1* with those of the buffer pointed to by *s2*.

### **Returns:**

0             The object pointed to by *s1* is the same as the object pointed to by *s2*.

Less than 0     The object pointed to by *s1* is less than the object pointed to by *s2*.

Greater than 0     The object pointed to by *s1* is greater than the object pointed to by *s2*.

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
```

```
char buffer[80];
int retval;

strcpy(buffer, "World");

retval = memcmp(buffer, "hello", 5);
if(retval < 0) {
 printf("Less than\n");
} else if(retval == 0) {
 printf("Equal\n");
} else {
 printf("Greater than\n");
}

return EXIT_SUCCESS;
}
```

produces the output:

```
Less than
```

## Classification:

QNX 4

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memcpy()*, *memchr()*, *memcmp()*, *memcpy()*, *memmove()*, *memset()*

## ***memmove()***

© 2004, QNX Software Systems Ltd.

*Copy bytes from one buffer to another, handling overlapping memory correctly*

### **Synopsis:**

```
#include <string.h>

void* memmove(void* dst,
 const void* src,
 size_t length);
```

### **Arguments:**

*dst*        A pointer to where you want the function to copy the data.

*src*        A pointer to the buffer that you want to copy data from.

*length*     The number of bytes to copy.

### **Library:**

libc

### **Description:**

The *memmove()* function copies *length* bytes from the buffer pointed to by *src* to the buffer pointed to by *dst*. Copying of overlapping regions is handled safely. Use *memcpy()* for greater speed when copying buffers that don't overlap.

### **Returns:**

A pointer to the destination buffer (that is, the value of *dst*).

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];

 strcpy(buffer, "World");
 memmove(buffer+1, buffer, 79);
```

```
 printf ("%s\n", buffer);
 return EXIT_SUCCESS;
}
```

produces the output:

```
WWorld
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memcpy()*, *memchr()*, *memcmp()*, *memcpy()*, *memcmp()*, *memset()*,  
*wmemmove()*

# ***memset()***

© 2004, QNX Software Systems Ltd.

*Set memory to a given value*

---

## **Synopsis:**

```
#include <string.h>

void* memset(void* dst,
 int c,
 size_t length);
```

## **Arguments:**

*dst*        A pointer to the memory that you want to set.

*c*         The value that you want to store in each byte.

*length*    The number of bytes to set.

## **Library:**

`libc`

## **Description:**

The *memset()* function fills *length* bytes starting at *dst* with the value *c*.

## **Returns:**

A pointer to the destination buffer (that is, the value of *dst*).

## **Examples:**

```
#include <string.h>
#include <stdlib.h>
#include <stdio.h>

int main(void)
{
 char buffer[80];

 memset(buffer, '=', 80);
 buffer[79] = '\0';

 puts(buffer);
}
```



```
 return EXIT_SUCCESS;
 }
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memccpy()*, *memchr()*, *memcmp()*, *memcpy()*, *memicmp()*, *memmove()*

## ***message\_attach()***

© 2004, QNX Software Systems Ltd.

*Attach a message range*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int message_attach(dispatch_t * dpp,
 message_attr_t * attr,
 int low,
 int high,
 int (* func) (
 message_context_t * ctx,
 int code,
 unsigned flags,
 void * handle),
 void * handle);
```

### **Arguments:**

|                  |                                                                                                                                                              |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>dpp</i>       | The dispatch handle, as returned by <i>dispatch_create()</i> .                                                                                               |
| <i>attr</i>      | A pointer to a <b>message_attr_t</b> structure that lets you specify additional requirements for the message; see “ <b>message_attr_t</b> structure,” below. |
| <i>low, high</i> | The range of messages that you’re interested in.                                                                                                             |
| <i>func</i>      | The function that you want to call when a message in the given range is received; see “Handler function,” below.                                             |
| <i>handle</i>    | An arbitrary handle that you want to associate with data for the defined message range. This handle is passed to <i>func</i> .                               |

### **Library:**

libc

## Description:

The *message\_attach()* function attaches a handler to the message range defined by the message type [*low*, *high*] (i.e. an inclusive message range) for dispatch handle *dpp*.



---

It's considered a programming error to attach overlapping message or pulse ranges. Message types should be greater than `IO_MAX` (defined in `<sys/iomsg.h>`).

---

When a message with a type in that range is received, *dispatch\_handler()* calls the user-supplied function *func*. You can also use the same function with *pulse\_attach()*. By examining *ctp->rvid*, *func* can determine whether a pulse or message was received.

This function is responsible for doing any specific work needed to handle the message pointed to by *ctp->msg*. The *handle* passed to the function is the *handle* initially passed to *message\_attach()*.

## message\_attr\_t structure

The *attr* argument is a pointer to a `message_attr_t` structure:

```
typedef struct _message_attr {
 unsigned flags;
 unsigned nparts_max;
 unsigned msg_max_size;
} message_attr_t;
```

You can use this structure to specify:

- the maximum message size to be received (the context allocated must be at least big enough to contain a message of that size)
- the maximum number of iovs to reserve in the `message_context_t` structure (*attr->nparts\_max*)
- various *flags*:

Currently, the following *attr->flags* are defined:

## MSG\_FLAG\_CROSS\_ENDIAN

Allow the server to receive messages from clients on machines with different native endian formats.

## MSG\_FLAG\_DEFAULT\_FUNC

Call this function if no other match is found, in this case, *low* and *high* are ignored. This overrides the default behavior of *dispatch\_handler()* which is to return *MsgError()* (ENOSYS) to the sender when an unknown message is received.

## Handler function

The user-supplied function *func* is called when a message in the defined range is received. This function is passed the message context *ctp*, in which the message was received, the message type, and the *handle* (the one passed to *message\_attach()*). Currently, the argument *flags* is reserved. Your function should return 0; other return values are reserved.

Here's a brief description of the context pointer fields:

|                      |                                               |
|----------------------|-----------------------------------------------|
| <i>ctp-&gt;rcvid</i> | The receive ID of the message.                |
| <i>ctp-&gt;msg</i>   | A pointer to the message.                     |
| <i>ctp-&gt;info</i>  | Data from a <code>_msg_info</code> structure. |

## Returns:

Zero on success, or -1 on failure (*errno* is set).

## Errors:

|        |                                             |
|--------|---------------------------------------------|
| EINVAL | The message <i>code</i> is out of range.    |
| ENOMEM | Insufficient memory to attach message type. |

**Examples:**

In this example, we create a resource manager where we attach to a private message range and attach a pulse, which is then used as a timer event:

```
#include <stdio.h>
#include <stddef.h>
#include <stdlib.h>

#define THREAD_POOL_PARAM_T dispatch_context_t
#include <sys/iofunc.h>
#include <sys/dispatch.h>

static resmgr_connect_funcs_t connect_func;
static resmgr_io_funcs_t io_func;
static iofunc_attr_t attr;

int
timer_tick(message_context_t *ctp, int code,
 unsigned flags, void *handle) {

 union sigval value = ctp->msg->pulse.value;
 /* Do some useful work on every timer firing... */
 printf("received timer event, value %d\n", value.sival_int);
 return 0;
}

int
message_handler(message_context_t *ctp, int code,
 unsigned flags, void *handle) {
 printf("received private message, type %d\n", code);
 return 0;
}

int
main(int argc, char **argv) {
 thread_pool_attr_t pool_attr;
 thread_pool_t *tpp;
 dispatch_t *dpp;
 resmgr_attr_t resmgr_attr;
 int id;
 int timer_id;
 struct sigevent event;
 struct _itimer itime;

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr,
 "%s: Unable to allocate dispatch handle.\n",
```

```
 argv[0]);
 return EXIT_FAILURE;
}

memset(&pool_attr, 0, sizeof pool_attr);
pool_attr.handle = dpp;
/* We are doing resmgr and pulse-type attaches */
pool_attr.context_alloc = dispatch_context_alloc;
pool_attr.block_func = dispatch_block;
pool_attr.unblock_func = dispatch_unblock;
pool_attr.handler_func = dispatch_handler;
pool_attr.context_free = dispatch_context_free;
pool_attr.lo_water = 2;
pool_attr.hi_water = 4;
pool_attr.increment = 1;
pool_attr.maximum = 50;

if((tpp = thread_pool_create(&pool_attr,
 POOL_FLAG_EXIT_SELF)) == NULL) {
 fprintf(stderr,
 "%s: Unable to initialize thread pool.\n",
 argv[0]);
 return EXIT_FAILURE;
}

iofunc_func_init(_RESMGR_CONNECT_NFUNCS, &connect_func,
 _RESMGR_IO_NFUNCS, &io_func);
iofunc_attr_init(&attr, S_IFNAM | 0666, 0, 0);

memset(&resmgr_attr, 0, sizeof resmgr_attr);
resmgr_attr.nparts_max = 1;
resmgr_attr.msg_max_size = 2048;

if((id = resmgr_attach(dpp, &resmgr_attr, "/dev/mynull",
 _FTYPE_ANY, 0,
 &connect_func, &io_func, &attr)) == -1) {
 fprintf(stderr, "%s: Unable to attach name.\n", argv[0]);
 return EXIT_FAILURE;
}

/*
 * We want to handle our own private messages, of type
 * 0x5000 to 0x5fff
 */
if(message_attach(dpp, NULL, 0x5000, 0x5fff,
 &message_handler, NULL) == -1) {
 fprintf(stderr,
 "Unable to attach to private message range.\n");
 return EXIT_FAILURE;
}
}
```

```

/* Initialize an event structure, and attach a pulse to it */
if((event.sigev_code = pulse_attach(dpp,
 MSG_FLAG_ALLOC_PULSE, 0,
 &timer_tick, NULL)) == -1) {
 fprintf(stderr, "Unable to attach timer pulse.\n");
 return EXIT_FAILURE;
}

/* Connect to our channel */
if((event.sigev_coid = message_connect(dpp,
 MSG_FLAG_SIDE_CHANNEL)) == -1) {
 fprintf(stderr, "Unable to attach to channel.\n");
 return EXIT_FAILURE;
}

event.sigev_notify = SIGEV_PULSE;
event.sigev_priority = -1;
/*
 * We could create several timers and use different
 * sigev values for each
 */
event.sigev_value.sival_int = 0;

if((timer_id = TimerCreate(CLOCK_REALTIME, &event)) == -1) {;
 fprintf(stderr,
 "Unable to attach channel and connection.\n");
 return EXIT_FAILURE;
}

/* And now setup our timer to fire every second */
itime.nsec = 1000000000;
itime.interval_nsec = 1000000000;
TimerSettime(timer_id, 0, &itime, NULL);

/* Never returns */
thread_pool_start(tpp);
return EXIT_SUCCESS;
}

```

For more examples using the dispatch interface, see *dispatch\_create()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*dispatch\_block(), dispatch\_create(), dispatch\_handler(),  
dispatch\_unblock(), message\_connect(), message\_detach(),  
\_msg\_info, pulse\_attach()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.



### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int message_connect(dispatch_t * dpp,
 int flags);
```

### **Arguments:**

- dpp*      The dispatch handle, as returned by *dispatch\_create()*.
- flags*    Currently, the following flag is defined in *<sys/dispatch.h>*:
- `MSG_FLAG_SIDE_CHANNEL` — request the connection ID be returned from a different space. This ID will be greater than any valid file descriptor. Once created there's no difference in the use of the messaging primitives on these IDs.

### **Library:**

`libc`

### **Description:**

The *message\_connect()* function creates a connection to the channel used by dispatch handle *dpp*. This function calls the *ConnectAttach()* kernel call. To detach the connection ID, you can call *ConnectDetach()*.



The `message_connect()` function works only when the dispatch blocking type is receive, i.e. attaches were done for `resmgr`, `message`, or select “type” events. If no attaches were done yet, the `message_connect()` call fails, since dispatch can’t determine if receive or `sigwait` blocking will be used.

## Returns:

A connection ID used by the message primitives, or -1 if an error occurs (`errno` is set).

## Errors:

- EAGAIN All kernel connection objects are in use.
- EINVAL Dispatch `dpp` doesn’t have a channel.

## Examples:

```
#include <sys/dispatch.h>
#include <time.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 dispatch_t *dpp;
 int flags, coid, id;

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr,
 "%s: Unable to allocate dispatch context.\n",
 argv[0]);
 return EXIT_FAILURE;
 }

 id = resmgr_attach (...);

 :

 if ((coid = message_connect (dpp, flags)) == -1) {
 fprintf (stderr, "Failed to create connection \
 to channel used by dispatch.\n");
 return 1;
 }
}
```

```
 }
 /* else connection to channel used by dispatch is created */

 :
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

Dispatch *dpp* must block on messages.

## See also:

*ConnectAttach()*, *message\_attach()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

## ***message\_detach()***

© 2004, QNX Software Systems Ltd.

*Detach a message range*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int message_detach(dispatch_t * dpp,
 int low,
 int high,
 int flags);
```

### **Arguments:**

|                  |                                                                                                                                              |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <i>dpp</i>       | The dispatch handle, as returned by <i>dispatch_create()</i> .                                                                               |
| <i>low, high</i> | The range of messages that you want to detach the handler from. This range must be the same one that you passed to <i>message_attach()</i> . |
| <i>flags</i>     | Reserved.                                                                                                                                    |

### **Library:**

libc

### **Description:**

The *message\_detach()* function detaches the message type [*low, high*], for dispatch handle *dpp*, that was attached with *message\_attach()*.

### **Returns:**

Zero on success. If an error occurs, -1 or the following error constant:

|        |                                                                                                         |
|--------|---------------------------------------------------------------------------------------------------------|
| EINVAL | The range [ <i>low, high</i> ] doesn't match the range that you attached with <i>message_attach()</i> . |
|--------|---------------------------------------------------------------------------------------------------------|

**Examples:**

```
#include <sys/dispatch.h>
#include <time.h>
#include <stdio.h>
#include <stdlib.h>

int my_func(...) {
 ...
}

int main(int argc, char **argv) {
 dispatch_t *dpp;
 int lo=0x2000, hi=0x2fff, flags=0;

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr,
 "%s: Unable to allocate dispatch handle.\n",
 argv[0]);
 return EXIT_FAILURE;
 }
 ...

 if(message_attach(dpp, NULL, lo, hi,
 &my_func, NULL) == -1) {
 fprintf(stderr,
 "%s: Failed to attach message range.\n",
 argv[0]);
 return 1;
 }
 ...

 if (message_detach (dpp, lo, hi, flags) == -1) {
 fprintf (stderr,
 "Failed to detach message range from %d to %d.\n",
 lo, hi);
 return 1;
 }
 /* else message was detached */
 ...
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*message\_attach()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

## Synopsis:

```
#include <stdlib.h>

#define min(a,b) ...
```

## Arguments:

*a,b* The numbers that you want to get the lesser of.

## Library:

libc

## Description:

The *min()* function returns the lesser of two values.



---

The *min()* function is for C programs only. For C+ and C++ programs, use the *\_max()* or *\_min()* macros.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int a;

 a = min(1, 10);
 printf("The value is: %d\n", a);
 return EXIT_SUCCESS;
}
```

## Classification:

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*max()*



## Synopsis:

```
#include <sys/types.h>
#include <sys/stat.h>

int mkdir(const char *path,
 mode_t mode);
```

## Arguments:

- path*      The name of the directory that you want to create.
- mode*      The permissions for the directory, modified by the process's file-creation mask (see *umask()*).
- The access permissions for the file or directory are specified as a combination of bits defined in the `<sys/stat.h>` header file. For more information, see "Access permissions" in the documentation for *stat()*.

## Library:

`libc`

## Description:

The *mkdir()* function creates a new subdirectory named *path*. The *path* can be relative to the current working directory or it can be an absolute path name.

The directory's owner ID is set to the process's effective user ID. The directory's group ID is set to the group ID of the parent directory (if the parent set-group ID bit is set) or to the process's effective group ID.

The newly created directory is empty.

The *mkdir()* function marks the *st\_atime*, *st\_ctime*, and *st\_mtime* fields of the directory for update. Also, the *st\_ctime* and *st\_mtime* fields of the parent directory are also updated.

**Returns:**

0, or -1 if an error occurs (*errno* is set).

**Errors:**

|              |                                                                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of <i>path</i> , or write permission is denied on the parent directory of <i>path</i> . |
| EEXIST       | The directory named by <i>path</i> already exists.                                                                                  |
| ELOOP        | Too many levels of symbolic links.                                                                                                  |
| EMLINK       | The link count of the parent directory would exceed LINK_MAX.                                                                       |
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.                                        |
| ENOENT       | A pathname component in the specified <i>path</i> does not exist, or <i>path</i> is an empty string.                                |
| ENOSPC       | The filesystem does not contain enough space to hold the contents of the new directory or to extend the parent directory.           |
| ENOSYS       | This function is not supported for this path.                                                                                       |
| ENOTDIR      | A component of <i>path</i> is not a directory.                                                                                      |
| EROFS        | The parent directory resides on a read-only filesystem.                                                                             |

**Examples:**

To make a new directory called */src* in */hd*:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <stdlib.h>

int main(void)
```

```
{
 (void)mkdir("/hd/src",
 S_IRWXU |
 S_IRGRP | S_IXGRP |
 S_IROTH | S_IXOTH);

 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*chdir()*, *chmod()*, *errno*, *getcwd()*, *mknod()*, *rmdir()*, *stat()*, *umask()*

# **mkfifo()**

© 2004, QNX Software Systems Ltd.

Create a FIFO special file

---

## **Synopsis:**

```
#include <sys/types.h>
#include <sys/stat.h>

int mkfifo(const char* path,
 mode_t mode);
```

## **Arguments:**

*path*      The pathname that you want to use for the FIFO special file.

*mode*      The file permission bits for the new FIFO. For more information, see “Access permissions” in the documentation for *stat()*.

## **Library:**

libc

## **Description:**

The *mkfifo()* function creates a new FIFO special file named by the pathname pointed to by *path*. The file permission bits of the new FIFO are initialized from *mode*, modified by the process’s creation mask (see *umask()*). Bits that are set in *mode* other than the file permission bits are ignored.

The FIFO owner ID is set to the process’s effective user ID and the FIFO’s group ID is set to the process’s effective group ID.

If *mkfifo()* succeeds, the *st\_mtime*, *st\_ctime*, *st\_atime* and *st\_mtime* fields of the file are marked for update. Also, the *st\_ctime* and *st\_mtime* fields of the directory that contains the new entry are marked for update.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EACCES A component of the path prefix denies search permission.
- EEXIST The named file already exists.
- ENAMETOOLONG  
The length of the *path* string exceeds PATH\_MAX, or a pathname component is longer than NAME\_MAX.
- ENOENT A component of the path prefix doesn't exist, or the *path* arguments points to an empty string.
- ENOSPC The directory that would contain the new file cannot be extended, or the filesystem is out of file allocation resources (that is, the disk is full).
- ENOSYS This function isn't supported for this path.
- ENOTDIR A component of the path prefix isn't a directory.
- EROFS The named file resides on a read-only filesystem.

**Examples:**

```
#include <sys/types.h>
#include <sys/stat.h>
#include <stdlib.h>

int main(void)
{
 (void)mkfifo("hd/qnx", S_IRUSR | S_IWUSR);

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*chmod(), errno, mknod(), pipe(), stat(), umask()*

**Synopsis:**

```
#include <sys/types.h>
#include <unistd.h>
#include <sys/stat.h>

int mknod(const char * path,
 mode_t mode,
 dev_t dev);
```

**Arguments:**

- path*      The pathname that you want to use for the file.
- mode*      A set of bits that define the file type and access permissions that you want to use. The valid file types are:
- S\_IFDIR — create a directory.
  - S\_IFIFO — create a FIFO.
- For more information, see “Access permissions” in the documentation for *stat()*.
- dev*      Ignored.

**Library:**

`libc`

**Description:**

The *mknod()* makes a file, named *path*, using the file type encoded in the *mode* argument. Supported file types are directories and FIFOs.



---

This function is included to enhance portability with software written for Unix-compatible operating systems. For POSIX portability, use *mkdir()* or *mkfifo()* instead.

---

To make a directory with read-write-execute permissions for everyone, you could use the following:

```
mknod (name, S_IFDIR | 0777, 0);
```

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EACCES A component of the *path* prefix denies search permission, or write permission is denied for the parent directory.
- EEXIST The named file already exists.
- ELOOP Too many levels of symbolic links or prefixes.
- EMLINK The link count of the parent directory would exceed LINK\_MAX.
- ENAMETOOLONG  
The length of the *path* string exceeds PATH\_MAX, or a pathname component is longer than NAME\_MAX.
- ENOENT A component of the *path* prefix doesn't exist, or the *path* arguments points to an empty string.
- ENOSPC The directory that would contain the new file cannot be extended or the filesystem is out of file allocation resources (that is, the disk is full).
- ENOSYS The *mknod()* function isn't implemented for the filesystem specified in *path*.
- ENOTDIR A component of the *path* prefix isn't a directory.
- EROFS The named file resides on a read-only filesystem.



**Examples:**

```

/*
 * Create special files as a directory or FIFO
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>
#include <unistd.h>

int main(int argc, char** argv)
{
 int c;
 mode_t mode = 0666;
 int ecode = 0;

 if(argc == 1) {
 printf("Use: %s [-d directory] ... [-f fifo] ... \n",
 argv[0]);
 return(0);
 }

 while((c = getopt(argc, argv, "d:f:")) != -1) {
 switch(c) {
 case 'd': mode = S_IFDIR | 0666; break;
 case 'f': mode = S_IFIFO | 0666; break;
 }

 if(mknod(optarg, mode, 0) != 0) {
 perror(optarg);
 ++ecode;
 }
 }

 return(ecode);
}

```

**Classification:**

Standard Unix

**Safety**

Cancellation point Yes

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*errno, mkdir(), mkfifo()*

## Synopsis:

```
#include <stdlib.h>

int mkstemp(char* template);
```

## Arguments:

*template*      A template for the filename that you want to use. This template can be any file name with some number of Xs appended to it, for example `/tmp/temp.XXXX`.

## Library:

`libc`

## Description:

The *mkstemp()* function takes the given file name template and overwrites a portion of it to create a filename. This file name is unique and suitable for use by the application. The trailing Xs are replaced with the current process number and/or a unique letter combination. The number of unique file names *mkstemp()* can return depends on the number of Xs provided; if you specify six Xs, *mkstemp()* tests roughly  $26^6$  combinations.

The *mkstemp()* function (unlike *mktemp()*) creates the template file, mode 0600 (i.e. read-write for the owner), returning a file descriptor opened for reading and writing. This avoids the race between testing for a file's existence and opening it for use.

## Returns:

The file descriptor of the temporary file, or -1 if no suitable file could be created; *errno* is set.

**Errors:**

ENOTDIR     The pathname portion of the template isn't an existing directory.

This function may also set *errno* to any value specified by *open()* and *stat()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

It's possible to run out of letters. The *mkstemp()* function doesn't check to determine whether the file name part of template exceeds the maximum allowable filename length.

For portability with X/Open standards prior to XPG4v2, use *tmpfile()* instead.

**See also:**

*chmod()*, *getpid()*, *mktemp()*, *open()* *stat()*, *tmpfile()*, *tmpnam()*

## Synopsis:

```
#include <stdlib.h>

char* mktemp(char* template);
```

## Arguments:

*template*      A template for the filename that you want to use. This template can be any file name with some number of Xs appended to it, for example `/tmp/temp.XXXX`.

## Library:

`libc`

## Description:

The *mktemp()* function takes the given file name template and overwrites a portion of it to create a filename. This file name is unique and suitable for use by the application. The trailing Xs are replaced with the current process number and/or a unique letter combination. The number of unique file names *mktemp()* can return depends on the number of Xs provided; if you specify six Xs, *mktemp()* tests roughly  $26^6$  combinations.

The *mkstemp()* function (unlike this function) creates the template file, mode 0600 (i.e. read-write for the owner), returning a file descriptor opened for reading and writing. This avoids the race between testing for a file's existence and opening it for use.

## Returns:

A pointer to the template, or NULL on failure; *errno* is set.

## Errors:

ENOTDIR      The pathname portion of the template isn't an existing directory.

This function may also set *errno* to any value specified by *stat()*.

**Classification:**

Legacy Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

In general, avoid using *mktemp()*, because a hostile process can exploit a race condition in the time between the generation of a temporary filename by *mktemp()* and the invoker's use of the temporary name. Use *mkstemp()* instead.

This function can create only 26 unique file names per thread for each unique template.

**See also:**

*chmod()*, *getpid()*, *mkstemp()*, *open()*, *stat()*, *tmpfile()*, *tmpnam()*

## Synopsis:

```
#include <time.h>

time_t mktime(struct tm* timeptr);
```

## Arguments:

*timeptr*     A pointer to a **tm** structure that contains the local time that you want to convert.

## Library:

**libc**

## Description:

The *mktime()* function converts the local time information in the **struct tm** specified by *timeptr* into a calendar time (Coordinated Universal Time) with the same encoding used by the *time()* function.

The original values of the *tm\_sec*, *tm\_min*, *tm\_hour*, *tm\_mday* and *tm\_mon* fields aren't restricted to the ranges described for **struct tm**. If these fields aren't in their proper ranges, they're adjusted so that they are. Values for the fields *tm\_wday* and *tm\_yday* are computed after all the other fields have been adjusted.

The original value of *tm\_isdst* is interpreted as follows:

- < 0     This field is computed as well.
- 0     Daylight savings time isn't in effect.
- > 0     Daylight savings time is in effect.

Whenever *mktime()* is called, the *tzset()* function is also called.

**Returns:**

The converted calendar time, or -1 if *mktime()* can't convert it.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

static const char *week_day[] = {
 "Sunday", "Monday", "Tuesday", "Wednesday",
 "Thursday", "Friday", "Saturday"
};

int main(void)
{
 struct tm new_year;
 time_t t;

 new_year.tm_year = 2001 - 1900;
 new_year.tm_mon = 0;
 new_year.tm_mday = 1;
 new_year.tm_hour = 0;
 new_year.tm_min = 0;
 new_year.tm_sec = 0;
 new_year.tm_isdst = 0;

 t = mktime(&new_year);
 if (t == (time_t)-1)
 printf("No conversion possible.\n");
 else
 printf("The 21st century begins on a %s.\n",
 week_day[new_year.tm_wday]);

 return EXIT_SUCCESS;
}
```

produces the output:

```
The 21st century begins on a Monday.
```



## Classification:

ANSI, POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*asctime()*, *asctime\_r()*, *clock()*, *ctime()*, *ctime\_r()*, *difftime()*, *gmtime()*, *gmtime\_r()*, *localtime()*, *localtime\_r()*, *strftime()*, *time()*, **tm**, *tzset()*

## ***mlock()***

© 2004, QNX Software Systems Ltd.

*Lock a buffer in physical memory*

---

### **Synopsis:**

```
#include <sys/mman.h>

int mlock(const void * addr,
 size_t len);
```

### **Library:**

libc

### **Description:**

The *mlock()* function isn't currently supported.

### **Returns:**

-1 to indicate an error (*errno* is set).

### **Errors:**

ENOSYS      The *mlock()* function isn't currently supported.

### **Classification:**

POSIX 1003.1 (Realtime Extensions)

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mlockall(), munlock(), munlockall()*

## ***mlockall()***

© 2004, QNX Software Systems Ltd.

*Lock a process's address space*

---

### **Synopsis:**

```
#include <sys/mman.h>

int mlockall(int flags);
```

### **Library:**

libc

### **Description:**

The *mlockall()* function isn't currently supported.

### **Returns:**

-1 to indicate an error (*errno* is set.)

### **Errors:**

ENOSYS      The *mlockall()* function isn't currently supported.

### **Classification:**

POSIX 1003.1 (Realtime Extensions)

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mlock()*, *munlock()*, *munlockall()*

## ***mmap()*, *mmap64()***

© 2004, QNX Software Systems Ltd.

*Map a memory region into a process's address space*

### **Synopsis:**

```
#include <sys/mman.h>

void * mmap(void * addr,
 size_t len,
 int prot,
 int flags,
 int fildes,
 off_t off);

void * mmap64(void * addr,
 size_t len,
 int prot,
 int flags,
 int fildes,
 off64_t off);
```

### **Arguments:**

- addr*      NULL, or a pointer to where you want the object to be mapped in the calling process's address space.
- len*        The number of bytes to map into the caller's address space. It can't be 0.
- prot*       The access capabilities that you want to use for the memory region being mapped. You can combine at least the following protection bits, as defined in `<sys/mman.h>`:
- PROT\_EXEC — the region can be executed.
  - PROT\_NOCACHE — disable caching of the region (e.g. so it can be used to access dual-ported memory).
  - PROT\_NONE — the region can't be accessed.
  - PROT\_READ — the region can be read.
  - PROT\_WRITE — the region can be written.
- flags*      Flags that specify further information about handling the mapped region; see below.

- fdes* The file descriptor for a shared memory object, or NOFD if you're mapping physical memory.
- off* The offset into shared memory of the location that you want to start mapping.

## Library:

`libc`

## Description:

The *mmap()* function maps a region within the object beginning at *off* and continuing for *len* into the caller's address space and returns the location.

Typically, you don't need to use *addr*; you can just pass NULL instead. If you set *addr* to a non-NULL value, whether the object is mapped depends on whether or not you set MAP\_FIXED in *flags*:

MAP\_FIXED is set

The object is mapped to the address in *addr*, or the function fails.

MAP\_FIXED isn't set

The value of *addr* is taken as a hint as to where to map the object in the calling process's address space. The mapped area won't overlay any current mapped areas.

There are two parts to the *flags* parameter. The first part is a type (masked by the MAP\_TYPE bits), which you must specify as one of the following:

- MAP\_PRIVATE The mapping is private to the calling process. It allocates system RAM and copies the current object.
- MAP\_SHARED The mapping may be shared by many processes.

You can OR the following flags into the above type to further specify the mapping:

**MAP\_ANON** This is most commonly used with **MAP\_PRIVATE**. The *fildev* parameter must be **NOFD**. The allocated memory is zero-filled. This is equivalent to opening */dev/zero*.

**MAP\_BELOW16M** Used with **MAP\_PHYS** | **MAP\_ANON**. The allocated memory area resides in physical memory below 16M. This is important for using DMA with ISA bus devices.

**MAP\_FIXED** Map the object to the address specified by *addr*. If this area is already mapped, the call changes the existing mapping of the area.



---

Use **MAP\_FIXED** with caution. Not all memory models support it. In general, you should assume that you can **MAP\_FIXED** only at an address (and size) that a call to *mmap()* without **MAP\_FIXED** returned.

---

A memory area being mapped with **MAP\_FIXED** is first unmapped by the system using the same memory area. See *munmap()* for details.

**MAP\_LAZY** Delay acquiring system memory, and copying or zero-filling the **MAP\_PRIVATE** or **MAP\_ANON** pages, until an access to the area has occurred. If you set this flag, and there's no system memory at the time of the access, the thread gets a **SIGBUS** with a code of **BUS\_ADRERR**. This flag is a hint to the memory manager.

**MAP\_PHYS** Physical memory is required. The *fildev* parameter must be **NOFD**. When used with **MAP\_PRIVATE** or **MAP\_SHARED**, the offset specifies the exact



physical address to map (e.g. for video frame buffers), and is equivalent to opening `/dev/mem`. If used with `MAP_ANON`, then physically contiguous memory is allocated.

`MAP_NOX64K` and `MAP_BELOW16M` are used to further define the `MAP_ANON` allocated memory (useful on x86 only).




---

You should use `mmap_device_memory()` instead of `MAP_PHYS`.

---

`MAP_NOX64K` (Useful on x86 only). Used with `MAP_PHYS` | `MAP_ANON`. Prevent the allocated memory area from crossing a 64K boundary. This may be important to some DMA devices. If more than 64K is requested, the area begins on a 64K boundary.

`MAP_STACK` This flag tells the memory allocator what the `MAP_ANON` memory will be used for. It's only a hint.

Using the mapping flags described above, a process can easily share memory between processes:

```
/* Map in a shared memory region */
fd = shm_open("/datapoints", O_RDWR, 0777);
addr = mmap(0, len, PROT_READ|PROT_WRITE, MAP_SHARED, fd, 0);
```

To share memory with hardware such as video memory on an x86 platform:

```
/* Map in VGA display memory */
addr = mmap(0,
 65536,
 PROT_READ|PROT_WRITE,
 MAP_PHYS|MAP_SHARED,
 NOFD,
 0xa0000);
```

To allocate a DMA buffer for a bus-mastering PCI network card:

```
/* Allocate a physically contiguous buffer */
addr = mmap(0,
 262144,
 PROT_READ | PROT_WRITE | PROT_NOCACHE,
 MAP_PHYS | MAP_ANON,
 NOFD,
 0);
```

## Returns:

The address of the mapped-in object, or MAP\_FAILED if an error occurred (*errno* is set).

## Errors:

- |        |                                                                                                                                                                                                       |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES | The file descriptor in <i>fdes</i> isn't open for reading, or you specified PROT_WRITE and MAP_SHARED, and <i>fdes</i> isn't open for writing.                                                        |
| EBADF  | Invalid file descriptor, <i>fdes</i> .                                                                                                                                                                |
| EINVAL | Invalid <i>flags</i> type, or <i>len</i> is 0.                                                                                                                                                        |
| ENODEV | The <i>fdes</i> argument refers to an object for which <i>mmap()</i> is meaningless (e.g. a terminal).                                                                                                |
| ENOMEM | You specified MAP_FIXED, and the address range requested is outside of the allowed process address range, or there wasn't enough memory to satisfy the request.                                       |
| ENXIO  | The address from <i>off</i> for <i>len</i> bytes is invalid for the requested object, or you specified MAP_FIXED, and <i>addr</i> , <i>len</i> , and <i>off</i> are invalid for the requested object. |

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <fcntl.h>
#include <errno.h>
#include <stdlib.h>
```

```
#include <sys/mman.h>

int main(int argc, char *argv[])
{
 int i;
 unsigned char *addr, c;

 /* Map BIOS ROM */
 addr = mmap(0, 0x10000, PROT_READ | PROT_WRITE,
 MAP_SHARED | MAP_PHYS, NOFD, 0xf0000);
 if (addr == MAP_FAILED) {
 fprintf(stderr, "mmap failed : %s\n",
 strerror(errno));
 return EXIT_FAILURE;
 }
 printf("Map addr is %p\n", (void*) addr);

 for (i = 0; i < 3 * 80; ++i) {
 c = *addr++;
 if (c >= ' ' && c <= 0x7f)
 putchar(c);
 else
 putchar('.');
 }

 return EXIT_SUCCESS;
}
```

## Classification:

*mmap()* is POSIX 1003.1; *mmap64()* is for large-file support

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*mmap\_device\_memory()*, *munmap()*

## Synopsis:

```
#include <stdint.h>
#include <sys/mman.h>

uintptr_t mmap_device_io(size_t len,
 uint64_t io);
```

## Arguments:

*len*     The number of bytes of device I/O memory that you want to access. It can't be 0.

*io*     The address of the area that you want to access.

## Library:

`libc`

## Description:

The *mmap\_device\_io()* function maps *len* bytes of device I/O memory at *io* and makes it accessible via the *in\*()* and *out\*()* functions in `<hw/inout.h>`.

## Returns:

A handle to the device's I/O memory, or `MAP_DEVICE_FAILED` if an error occurs (*errno* is set).

## Errors:

`EINVAL`     Invalid *flags* type, or *len* is 0.

`ENOMEM`     The address range requested is outside of the allowed process address range, or there wasn't enough memory to satisfy the request.

`ENXIO`     The address from *io* for *len* bytes is invalid.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

You need I/O privileges to use the result of the *mmap\_device\_io()* function. The calling thread may call *ThreadCtl()* with the `_NTO_TCTL_IO` command to establish these privileges.

**See also:**

*mmap()*, *mmap\_device\_memory()*, *munmap\_device\_io()*

**Synopsis:**

```
#include <sys/mman.h>

void * mmap_device_memory(void * addr,
 size_t len,
 int prot,
 int flags,
 uint64_t physical);
```

**Arguments:**


*addr* NULL, or a pointer to where you want to map the object in the calling process's address space.

*len* The number of bytes you want to map into the caller's address space. It can't be 0.

*prot* The access capabilities that you want to use for the memory region being mapped. You can use a combination of at least the following protection bits, as defined in `<sys/mman.h>`:

- PROT\_EXEC — the region can be executed.
- PROT\_NOCACHE — disable the caching of the region (e.g. to access dual-ported memory).

---

 Read the architecture guide for your processor; you may need to add special instructions. For example, if you specify PROT\_NOCACHE on a PPC device, you may need to issue special *eieio()* instructions to ensure that writes occur in a desired order.

---

- PROT\_NONE — the region can't be accessed.
- PROT\_READ — the region can be read.
- PROT\_WRITE — the region can be written.

*flags* Specifies further information about handling the mapped region. You can use the following flag:

- MAP\_FIXED — map the object to the address specified by *addr*. If this area is already mapped, the call changes the existing mapping of the area.

☞ Use MAP\_FIXED with caution. Not all memory models support it. In general, you should assume that you can MAP\_FIXED only at an address (and size) that a call to *mmap()* without MAP\_FIXED returned. These restrictions will be removed from the user-to-user protection model with the full VM (available with a later version of the QNX Neutrino OS).

---

A memory area being mapped with MAP\_FIXED is first unmapped by the system using the same memory area. See *munmap()* for details.

This function already uses MAP\_SHARED ORed with MAP\_PHYS (see *mmap()* for a description of these flags).

*physical* The physical address of the memory to map into the caller's address space.

## Library:

libc

## Description:

The *mmap\_device\_memory()* function maps *len* bytes of a device's *physical* memory address into the caller's address space at the location returned by *mmap\_device\_memory()*.

You should use this function instead of using *mmap()* with the MAP\_PHYS flag.

Typically, you don't need to use *addr*; you can just pass NULL instead. If you set *addr* to a non-NULL value, whether the object is mapped depends on whether or not you set MAP\_FIXED in *flags*:

MAP\_FIXED is set

The object is mapped to the address in *addr*, or the function fails.



MAP\_FIXED isn't set

The value of *addr* is taken as a hint as to where to map the object in the calling process's address space. The mapped area won't overlay any current mapped areas.

## Returns:

The address of the mapped-in object, or MAP\_FAILED if an error occurs (*errno* is set).

## Errors:

- |        |                                                                                                                                                                                                                |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EINVAL | Invalid <i>flags</i> type, or <i>len</i> is 0.                                                                                                                                                                 |
| ENOMEM | The address range requested is outside of the allowed process address range, or there wasn't enough memory to satisfy the request.                                                                             |
| ENXIO  | The address from <i>physical</i> for <i>len</i> bytes is invalid for the requested object, or MAP_FIXED was specified and <i>addr</i> , <i>len</i> , and <i>physical</i> are invalid for the requested object. |

## Examples:

```
/* map in the physical memory, 0xb8000 is text mode VGA video memory */
ptr = mmap_device_memory(0, len, PROT_READ|PROT_WRITE|PROT_NOCACHE, 0, 0xb8000);
if (ptr == MAP_FAILED) {
 perror("mmap_device_memory for physical address 0xb8000 failed");
 exit(EXIT_FAILURE);
}
```

## Classification:

QNX Neutrino

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **Caveats:**

You need I/O privileges to use the result of the *mmap\_device\_memory()* function. The calling thread may call *ThreadCtl()* with the `_NTO_TCTL_IO` command to establish these privileges.

## **See also:**

*mmap()*, *mmap\_device\_io()*, *munmap\_device\_memory()*

## Synopsis:

```
#include <sys/modem.h>

int modem_open(char* device,
 speed_t baud);
```

## Arguments:

*device*     The path name of the serial port that you want to open.

*baud*       Zero, or the baud rate that you want to use.

## Library:

libc

## Description:

The *modem\_open()* function opens a serial port identified by *device*. The device is set to raw mode by changing the control flags using *tcgetattr()* and *tcsetattr()* as follows:

```
termio.c_cflag = CS8 | IHFLOW | OHFLOW | CREAD | HUPCL;
termio.c_iflag = BRKINT;
termio.c_lflag = IEXTEN;
termio.c_oflag = 0;
```

Any pending input or output characters are discarded.

If *baud* is nonzero, then the baud rate is changed to that value.

## Returns:

An open file descriptor, or -1 on failure (*errno* is set).

## Errors:

EACCES       Search permission is denied on a component of the path prefix, or the file doesn't exist.

|              |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADFSYS     | While attempting to open the named file, either the file itself or a component of the path prefix was found to be corrupted. A system failure — from which no automatic recovery is possible — occurred while the file was being written to, or while the directory was being updated. You'll need to invoke appropriate systems-administration procedures to correct this situation before proceeding. |
| EBUSY        | The file named by <i>device</i> is a block special device that's already open for writing, or <i>device</i> names a file that's on a filesystem mounted on a block special device that's already open for writing, or <i>device</i> is in use.                                                                                                                                                          |
| EINTR        | The open operation was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                                         |
| EISDIR       | The named <i>device</i> is a directory.                                                                                                                                                                                                                                                                                                                                                                 |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                                                                                                                                                                                                                                          |
| EMFILE       | Too many file descriptors are currently in use by this process.                                                                                                                                                                                                                                                                                                                                         |
| ENAMETOOLONG | The length of the <i>device</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.                                                                                                                                                                                                                                                                                               |
| ENFILE       | Too many files are currently open in the system.                                                                                                                                                                                                                                                                                                                                                        |
| ENOENT       | The named <i>device</i> doesn't exist, or the <i>path</i> argument points to an empty string.                                                                                                                                                                                                                                                                                                           |
| ENOSYS       | The <i>modem_open()</i> function isn't implemented for the filesystem specified in <i>device</i> .                                                                                                                                                                                                                                                                                                      |
| ENOTDIR      | A component of the path prefix isn't a directory.                                                                                                                                                                                                                                                                                                                                                       |
| ENXIO        | No process has the file open for reading.                                                                                                                                                                                                                                                                                                                                                               |

**Examples:**

```

#include <stdlib.h>
#include <unistd.h>
#include <string.h>
#include <fcntl.h>
#include <sys/modem.h>
#include <stdio.h>
#include <errno.h>

/*
 curstate curflags newstate newflags newtimeout
 newquiet retvalue pattern response
*/

struct modem_script table[] ={
 {1, 0, 1, 0, 2, 5, 0,
 NULL, "ATZ\\r\\P0a"},
 {1, 0, 2, 0, 30, 5, 0,
 "**ok*", "ATDT5910934"},
 {2, MODEM_BAUD, 3, MODEM_LASTLINE, 10, 5, 0,
 "**connect*", NULL},
 {3, 0, 4, 0, 8, 5, 0,
 "**login:*", "guest"},
 {4, MODEM_NOECHO, 5, 0, 15, 5, 0,
 "**password:*", "xxxx"},
 {5, 0, 0, 0, 0, 0, 0,
 "**$ *", NULL},
 {0, 0, 0, 0, 0, 0, 1,
 "**no carrier*", NULL},
 {0, 0, 0, 0, 0, 0, 2,
 "**no answer*", NULL},
 {0, 0, 0, 0, 0, 0, 3,
 "**no dialtone*", NULL},
 {0, 0, 0, 0, 0, 0, 4,
 "**busy**", NULL},
 { NULL }
};

void io(char* progress, char* in, char* out) {
 if(progress)
 printf("progress: %s\n", progress);

 if(in)
 printf("input: %s\n", in);

 if(out)
 printf("output: %s\n", out);
}

```

```
int main(int argc, char* argv[]) {
 int fd, status;
 speed_t baud = -1;

 if((fd = modem_open(argv[1], 0)) == -1) {
 fprintf(stderr, "Unable to open %s: %s\n",
 argv[1], strerror(errno));
 exit(1);
 }

 status = modem_script(fd, table, &baud,
 &io, NULL);
 printf("status=%d baud=%d\n", status, baud);
 exit(status);
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*modem\_read()*, *modem\_script()*, *modem\_write()*

**Synopsis:**

```
#include <sys/modem.h>

int modem_read(int fd,
 char* buf,
 int bufsize,
 int quiet,
 int timeout,
 int flags,
 int (*cancel) (void));
```

**Arguments:**

- fd* The file descriptor for the device that you want to read from; see *modem\_open()*.
- buf* A pointer to a buffer where the function can store the data.
- bufsize* The size of the buffer, in bytes.
- quiet* The maximum time to wait for more input after receiving at least one characters, in tenths of a second.
- timeout* The maximum time to wait for any input, in tenths of a second.
- flags* Flags that you can use to filter and map received characters; any combination of:
- MODEM\_ALLOWCASE — preserve the case of incoming characters. Without this flag, all letters are mapped to lower case.
  - MODEM\_ALLOWCTRL — allow control characters. Without this flag, control characters are discarded.
  - MODEM\_ALLOW8BIT — preserve the top bit of incoming characters. Without this flag, the top bit is set to zero for all characters.

- `MODEM_LASTLINE` — discard all previously received characters when a newline is received followed by more characters. Without this flag, *buf* may contain multiple lines. If an automatic login script may be presented with an arbitrary text screen before the login prompt, you can use this flag to discard all but the login line, reducing the possibility of false matches.

*cancel* NULL, or a callback that's called whenever the *quiet* time period expires while waiting for more input.

## Library:

`libc`

## Description:

The *modem\_read()* function reads up to *bufsize* bytes from the device specified by the file descriptor, *fd*, and places them into the buffer pointed to by *buf*.

If no characters are received within the given *timeout*, *modem\_read()* returns with -1.

When at least one character has been received, *modem\_read()* returns if the flow of incoming characters stops for at least the *quiet* time period. The number of characters saved in *buf* is returned.

If you provide a *cancel* function, it's called once each *quiet* time period while waiting for input. If this function returns a nonzero value, *modem\_read()* returns -1 immediately and sets *errno* to `ETIMEDOUT`. You can use the *cancel* function as a callback in a graphical dialer that needs to support a cancel button to stop a script (see *modem\_script()*).

## Returns:

Zero for success, or -1 on failure (*errno* is set).



**Errors:**

|        |                                                                                                                             |
|--------|-----------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set on this <i>fd</i> , and the process would have been blocked in trying to perform this operation. |
| EBADF  | The argument <i>fd</i> is invalid, or the file isn't opened for reading.                                                    |
| EINTR  | The <i>readcond()</i> call was interrupted by the process being signalled.                                                  |
| EIO    | This process isn't currently able to read data from this <i>fd</i> .                                                        |
| ENOSYS | This function isn't supported for this <i>fd</i> .                                                                          |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Yes                     |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Read the <i>Caveats</i> |

**Caveats:**

Depending on what you do in your *cancel* function, *modem\_read()* may or not be signal handler or thread-safe.

**See also:**

*modem\_open()*, *modem\_script()*, *modem\_write()*

## ***modem\_script()***

© 2004, QNX Software Systems Ltd.

*Run a script on a device*

### **Synopsis:**

```
#include <sys/modem.h>

int modem_script(int fd,
 struct modem_script* table,
 speed_t* baud,
 void (*io) (
 char* progress,
 char* in,
 char* out),
 int (*cancel) (void));
```

### **Arguments:**

|               |                                                                                                                                                |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>fd</i>     | The file descriptor for the device that you want to read from; see <i>modem_open()</i> .                                                       |
| <i>table</i>  | An array of <b>modem_script</b> structures that comprise a script of commands that you want to run on the device; see below.                   |
| <i>baud</i>   | A pointer to a <b>speed_t</b> where the function can store the baud rate (if you script says to do so).                                        |
| <i>io</i>     | A function that's called to process each string that's emitted or received.                                                                    |
| <i>cancel</i> | NULL, or a callback function that's called whenever the <i>newquiet</i> time period (specified in the script) expires while waiting for input. |

### **Library:**

**libc**

**Description:**

The *modem\_script()* function runs the script *table* on the device associated with the file descriptor *fd*. The script implements a simple state machine that emits strings and waits for responses.

Each string that's emitted or received is passed to the function *io()* as follows:

| Call                    | Description                    |
|-------------------------|--------------------------------|
| <i>(*io)(str, 0, 0)</i> | Emitted <i>progress</i> string |
| <i>(*io)(0, str, 0)</i> | Received string                |
| <i>(*io)(0, 0, str)</i> | Emitted <i>response</i> string |

This lets an application set up a callback that can display the script's interaction in a status window.

If you provide a *cancel* function, it's called once each *newquiet* 1/10 of a second while waiting for input. If this function returns a nonzero value, the read returns immediately with -1 and *errno* is set to ETIMEDOUT. You can use the *cancel* function as a callback in a graphical dialer that needs to support a cancel button to stop a script.

The *table* is an array of **modem\_script** structures that contain the following members:

|                             |                                                                                                                                                                                                                                                                          |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>char</b> <i>curstate</i> | The current state. Execution always begins at state 1, which must be the first array element of <i>table</i> . Multiple elements may have the same current state, in which case any received input is matched against each <i>pattern</i> member for that state.         |
| <b>int</b> <i>curflags</i>  | The flags to use on a pattern match of a response: <ul style="list-style-type: none"> <li>• MODEM_NOECHO — don't echo the response through the <i>io()</i> callback.</li> <li>• MODEM_BAUD — extract any number in the response and assign it to <i>baud</i>.</li> </ul> |

- char** *newstate* When a pattern match occurs with *pattern*, this is the next state. A state transition causes *response* to be output and *newflags*, *newtimeout*, and *newquiet* to be saved and associated with the new state. Changing to a new state of 0 causes *modem\_script()* to return with the value in *retvalue*.
- int** *newflags* Saved on a state transition and passed to *modem\_read()* when waiting for a response in the new state. For information about these flags, see *modem\_read()*.
- int** *newtimeout* Saved on a state transition and passed to *modem\_read()* when waiting for a response in the new state. This timeout is described in *modem\_read()*.
- int** *newquiet* Saved on a state transition and passed to *modem\_read()* when waiting for a response in the new state. This quiet timeout is described in *modem\_read()*.
- short** *retvalue* The return value when the script terminates with a pattern match, and the new state is 0.
- char\*** *pattern* A pattern to match against received characters. The pattern is matched using *fnmatch()*. Only patterns in the current state or the wildcard state of 0 are matched. On a match, the current state changes to *newstate*.
- char\*** *response* On a *pattern* match, this response is output to the device. If the *curflags* don't have `MODEM_NOECHO` set, the response is given to the callback function passed as the *io* parameter.
- char\*** *progress* On a *pattern* match, this progress string is passed to the callback function passed as the *io* parameter.

Here's an example that demonstrates the operation of the script:

```

/*
 curstate curflags newstate newflags newtimeout
 newquiet retvalue pattern response
*/

struct modem_script table[] ={
 {1, 0, 1, 0, 2, 5, 0,
 NULL, "ATZ\\r\\P0a"},
 {1, 0, 2, 0, 30, 5, 0,
 "*ok*", "ATDT5910934"},
 {2, MODEM_BAUD, 3, MODEM_LASTLINE, 10, 5, 0,
 "**connect*", NULL},
 {3, 0, 4, 0, 8, 5, 0,
 "**login:*", "guest"},
 {4, MODEM_NOECHO, 5, 0, 15, 5, 0,
 "**password:*", "xxxx"},
 {5, 0, 0, 0, 0, 0, 0,
 "**$ *", NULL},
 {0, 0, 0, 0, 0, 0, 1,
 "**no carrier*", NULL},
 {0, 0, 0, 0, 0, 0, 2,
 "**no answer*", NULL},
 {0, 0, 0, 0, 0, 0, 3,
 "**no dialtone*", NULL},
 {0, 0, 0, 0, 0, 0, 4,
 "**busy*", NULL},
 { NULL }
};

```

When this script is passed to *modem\_script()*, the current state is set to 1, and the output is **ATZ** (the *response* in the first array element).

While in any state, *modem\_script()* waits for input, matching it against the current state or the wildcard state of 0.

## State 1

| <b>Input</b>         | <b>Action</b>                                                                                                                                                                                                              |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>*ok*</b>          | Go to state 2 and emit <b>ATDT1-591-0934</b> . The flags to be used in the new state are set to 0, the quiet time in the new state is set to 5/10 of a second, and the timeout time in the new state is set to 30 seconds. |
| <b>*no carrier*</b>  | Go to state 0 (the termination newstate), return with the contents of <i>retvalue</i> (1).                                                                                                                                 |
| <b>*no answer*</b>   | Go to state 0 (the termination newstate), return with the contents of <i>retvalue</i> (2).                                                                                                                                 |
| <b>*no dialtone*</b> | Go to state 0 (the termination newstate), return with the contents of <i>retvalue</i> (3).                                                                                                                                 |
| <b>*busy*</b>        | Go to state 0 (the termination newstate), return with the contents of <i>retvalue</i> (4).                                                                                                                                 |

**State 2**

| <b>Input</b>         | <b>Action</b>                                                                                                                                                                                                                                                                                                                              |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>*connect*</b>     | Go to state 3 and don't emit anything to the device. The flags to be used in the new state are set to MODEM_LASTLINE, the quiet time in the new state is set to 5/10 of a second, and the timeout time in the new state is set to 10 seconds. Since the current flags are MODEM_BAUD, the baud rate is extracted from the connect message. |
| <b>*no carrier*</b>  | Same as previous table                                                                                                                                                                                                                                                                                                                     |
| <b>*no answer*</b>   | Same as previous table                                                                                                                                                                                                                                                                                                                     |
| <b>*no dialtone*</b> | Same as previous table                                                                                                                                                                                                                                                                                                                     |
| <b>*busy*</b>        | Same as previous table                                                                                                                                                                                                                                                                                                                     |

**State 3**

| <b>Input</b>         | <b>Action</b>                                                                                                                                                                                                    |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>*login*</b>       | Go to state 4 and emit <b>guest</b> . The flags to be used in the new state are set to 0, the quiet time in the new state is set to 5/10 of a second, and the timeout time in the new state is set to 8 seconds. |
| <b>*no carrier*</b>  | Same as previous table                                                                                                                                                                                           |
| <b>*no answer*</b>   | Same as previous table                                                                                                                                                                                           |
| <b>*no dialtone*</b> | Same as previous table                                                                                                                                                                                           |
| <b>*busy*</b>        | Same as previous table                                                                                                                                                                                           |

**State 4**

| <b>Input</b>         | <b>Action</b>                                                                                                                                                                                                                                                                                                                        |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>*password*</b>    | Go to state 5 and emit <b>xxxx</b> . The flags to be used in the new state are set to 0, the quiet time in the new state is set to 5/10 of a second, and the timeout time in the new state is set to 15 seconds. Since the current flags are MODEM_NOECHO, the password response <b>xxxx</b> isn't sent to the <i>io()</i> callback. |
| <b>*no carrier*</b>  | Same as previous table                                                                                                                                                                                                                                                                                                               |
| <b>*no answer*</b>   | Same as previous table                                                                                                                                                                                                                                                                                                               |
| <b>*no dialtone*</b> | Same as previous table                                                                                                                                                                                                                                                                                                               |
| <b>*busy*</b>        | Same as previous table                                                                                                                                                                                                                                                                                                               |

**State 5**

| Input         | Action                                                                                     |
|---------------|--------------------------------------------------------------------------------------------|
| *\$ *         | Go to state 0 (the termination newstate), return with the contents of <i>retvalue</i> (0). |
| *no carrier*  | Same as previous table                                                                     |
| *no answer*   | Same as previous table                                                                     |
| *no dialtone* | Same as previous table                                                                     |
| *busy*        | Same as previous table                                                                     |

If you set the flag MODEM\_BAUD for a state, then any number embedded in a matching response is extracted and assigned as a number to the *baud* parameter.

If you don't set the flag MODEM\_NOECHO for a state, then all emitted strings are also given to the passed *io* function as *(\*io) (0, 0, response)*.

### Returns:

The *retvalue* member of a script entry that terminates the script. This will always be a positive number. If *modem\_script* fails, it returns -1 and sets *errno*.

### Errors:

|        |                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the write operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for writing.                                                                                       |
| EINTR  | The write operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EIO    | A physical I/O error occurred. The precise meaning depends on the device.                                                                                                |



EPIPE      An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.

## Classification:

QNX Neutrino

### Safety

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Yes                     |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Read the <i>Caveats</i> |

## Caveats:

Depending on what you do in your *cancel* function, it might or might not be safe to call *modem\_script()* from a signal handler or a multithreaded program.

## See also:

*modem\_open()*, *modem\_read()*, *modem\_write()*

## ***modem\_write()***

© 2004, QNX Software Systems Ltd.

*Write a string to a device*

### **Synopsis:**

```
#include <sys/modem.h>

int modem_write(int fd,
 char* str);
```

### **Arguments:**

*fd*     The file descriptor for the device that you want to write to; see *modem\_open()*.

*str*     The string that you want to write.

### **Library:**

`libc`

### **Description:**

The *modem\_write()* function writes the string *str* to the device specified by the file descriptor *fd*. Just before writing each character, all buffered input from the same device is flushed. After writing each character, an attempt to read an echo is made. The intent is to write a string without its appearing back in the input stream even if the device is echoing each character written.

If the `\` character appears in *str*, then the character following it is interpreted by *modem\_write()*, and instead of both being written, they're treated as a special escape sequence that causes the following actions to be taken:

| <b>Escape</b>   | <b>Description</b>        |
|-----------------|---------------------------|
| <code>\r</code> | Output a carriage return. |
| <code>\n</code> | Output a newline.         |

*continued...*

| Escape            | Description                                                                 |
|-------------------|-----------------------------------------------------------------------------|
| <code>\xhh</code> | Output the single character whose hex representation follows as <i>hh</i> . |
| <code>\B</code>   | Send a 500 msec break on the line using <i>tcsendbreak()</i> .              |
| <code>\D</code>   | Drop the line for 1 second using <i>tcdropline()</i> .                      |
| <code>\Phh</code> | Pause for <i>hh</i> 1/10 of a second where <i>hh</i> is two hex characters. |

**Returns:**

Zero on success, -1 on failure (*errno* is set ).

**Errors:**

|        |                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the write operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for writing.                                                                                       |
| EINTR  | The write operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EIO    | A physical I/O error occurred. The precise meaning depends on the device.                                                                                                |
| EPIPE  | An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.                               |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*modem\_open(), modem\_read(), modem\_script()*

### **Synopsis:**

```
#include <math.h>

double modf(double value,
 double* iptr);

float modff(float value,
 float* iptr);
```

### **Arguments:**

*value*     The value that you want to break into parts.

*iptr*     A pointer to a location where the function can store the integral part of the number.

### **Library:**

**libm**

### **Description:**

The *modf()* and *modff()* functions break the given *value* into integral and fractional parts, each of which has the same sign as the argument. They store the integral part as a **double** in the object pointed to by *iptr*.

### **Returns:**

The signed fractional part of *value*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
{
 double integral_value, fractional_part;

 fractional_part = modf(4.5, &integral_value);
 printf("%f %f\n", fractional_part, integral_value);

 fractional_part = modf(-4.5, &integral_value);
 printf("%f %f\n", fractional_part, integral_value);

 return EXIT_SUCCESS;
}
```

produces the output:

```
0.500000 4.000000
-0.500000 -4.000000
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*frexp()*, *ldexp()*

**Synopsis:**

```
#include <sys/mount.h>

int mount(const char* spec,
 const char* dir,
 int flags,
 const char* type,
 const void* data,
 int datalen);
```

**Arguments:**

- spec*      A null-terminated string describing a special device (e.g. `/dev/hd0t77`), or NULL if there's no special device.
- dir*        A null-terminated string that names the directory that you want to mount (e.g. `/mnt/home`).
- flags*      Flags that are passed to the driver:
- `_MFLAG_OCB` — ignore the special device string, and contact all servers.
  - `_MOUNT_READONLY` — mark the filesystem mountpoint as read-only.
  - `_MOUNT_NOEXEC` — don't allow executables to load.
  - `_MOUNT_NOSUID` — don't honor setuid bits on the filesystem.
  - `_MOUNT_NOCREAT` — don't allow file creation on the filesystem.
  - `_MOUNT_OFF32` — limit `off_t` to 32 bits.
  - `_MOUNT_NOATIME` — disable logging of file access times.
  - `_MOUNT_BEFORE` — call `resmgr_attach()` with `RESMGR_FLAG_BEFORE`.

- `_MOUNT_AFTER` — call `resmgr_attach()` with `RESMGR_FLAG_AFTER`.
- `_MOUNT_OPAQUE` — call `resmgr_attach()` with `RESMGR_FLAG_OPAQUE`.
- `_MOUNT_UNMOUNT` — unmount this path.
- `_MOUNT_REMOUNT` — this path is already mounted; perform an update.
- `_MOUNT_FORCE` — force an unmount or a remount change.
- `_MOUNT_ENUMERATE` — autodetect on this device.

|                |                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>type</i>    | A null-terminated string with the filesystem type (e.g. <code>nfs</code> , <code>cifs</code> , <code>qnx4</code> , <code>ext2</code> , <code>network</code> ). |
| <i>data</i>    | A pointer to additional data to be sent to the manager. If <i>datalen</i> is <code>&lt;0</code> , the data points to a null-terminated string.                 |
| <i>datalen</i> | The length of the data, in bytes, that's being sent to the server, or <code>&lt;0</code> if the data is a null-terminated string.                              |

## Library:

`libc`

## Description:

The `mount()` function sends a request to servers to mount the services provided by *spec* and *type* at *dir*.

If you set `_MFLAG_OCB` in the flags, then the special device string is ignored, and all servers are contacted. If you don't set this bit, and the special device *spec* exists, then only the server that created that device is contacted, and the full path to *spec* is provided.

If *datalen* is any value `<0`, and there's a data pointer, the function assumes that the data pointer is a pointer to a string.



**Returns:**

-1 on failure; no server supports the request (*errno* is set).

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*resmgr\_attach()*, *umount()*

Writing a Resource Manager in *Programmer's Guide*

## ***mount\_parse\_generic\_args()***

© 2004, QNX Software Systems Ltd.

*Strip off common mount arguments*

### **Synopsis:**

```
#include <sys/mount.h>

char * mount_parse_generic_args(char * options,
 int * flags);
```

### **Arguments:**

*options*     The string of options that you want to parse; see below.

*flags*        A pointer to a location where the function can store a set of bits corresponding to the options that it finds; see below.

### **Library:**

libc

### **Description:**

The *mount\_parse\_generic\_args()* function allows you to strip out common flags to help you parse mount arguments. This is useful when you want to create a custom mount utility.

Here's a list of the supported options that may be stripped:

| <b>Option:</b> | <b>Set/Clear this bit:</b>      | <b>Description:</b>                                                |
|----------------|---------------------------------|--------------------------------------------------------------------|
| <b>after</b>   | Set <code>_MOUNT_AFTER</code>   | Call <i>resmgr_attach()</i> with <code>RESMGR_FLAG_AFTER</code> .  |
| <b>atime</b>   | Clear <code>_MOUNT_ETIME</code> | Log file access times (default).                                   |
| <b>before</b>  | Set <code>_MOUNT_BEFORE</code>  | Call <i>resmgr_attach()</i> with <code>RESMGR_FLAG_BEFORE</code> . |
| <b>creat</b>   | Clear <code>_MOUNT_CREAT</code> | Allow file creation on the filesystem (default).                   |

*continued...*

| <b>Option:</b>   | <b>Set/Clear this bit:</b>         | <b>Description:</b>                                                                        |
|------------------|------------------------------------|--------------------------------------------------------------------------------------------|
| <b>enumerate</b> | Set <code>_MOUNT_ENUMERATE</code>  | Auto-detect on this device.                                                                |
| <b>exec</b>      | Clear <code>_MOUNT_NOEXEC</code>   | Load executables (default).                                                                |
| <b>force</b>     | Set <code>_MOUNT_FORCE</code>      | Force an unmount or a remount change.                                                      |
| <b>noatime</b>   | Set <code>_MOUNT_NOATIME</code>    | Disable logging of file access times.                                                      |
| <b>nocreat</b>   | Set <code>_MOUNT_NOCREAT</code>    | Don't allow file creation on the filesystem.                                               |
| <b>noexec</b>    | Set <code>_MOUNT_NOEXEC</code>     | Don't allow executables to load.                                                           |
| <b>nostat</b>    | Set <code>_MFLAG_OCB</code>        | Don't attempt to <i>stat()</i> the special device before mounting (i.e. <code>-t</code> ). |
| <b>nosuid</b>    | Set <code>_MOUNT_NOSUID</code>     | Don't honor setuid bits on the filesystem.                                                 |
| <b>opaque</b>    | Set <code>_MOUNT_OPAQUE</code>     | Call <i>resmgr_attach()</i> with <code>RESMGR_FLAG_OPAQUE</code> .                         |
| <b>remount</b>   | Set <code>_MOUNT_REMOUNT</code>    | This path is already mounted; perform an update.                                           |
| <b>ro</b>        | Set <code>_MOUNT_READONLY</code>   | Mark the filesystem mountpoint as read-only.                                               |
| <b>rw</b>        | Clear <code>_MOUNT_READONLY</code> | Mark the filesystem mountpoint as read/write (default).                                    |
| <b>suid</b>      | Clear <code>_MOUNT_SUID</code>     | Honor setuid bits on the filesystem (default).                                             |
| <b>update</b>    | Set <code>_MOUNT_REMOUNT</code>    | This path is already mounted, perform an update.                                           |

## Returns:

A string pointing to unprocessed options.

## Examples:

```
while ((c = getopt(argv, argc, "o:")) {
 switch (c) {
 case 'o':
 if ((mysteryop = mount_parse_generic_args(optarg, &flags)) {
 /*
 * You can do your own getsubopt type processing here
 * mysteryop is stripped of the common options.
 */
 }
 break;
 }
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*mount()*, *resmgr\_attach()*, *umount()*

**mount** in the *Utilities Reference*

Writing a Resource Manager in *Programmer's Guide*

## Synopsis:

```
#include <malloc.h>

enum mcheck_status mprobe(void * ptr);
```

## Arguments:

*ptr*     A pointer to the start of the heap block.

## Library:

`libc`

## Description:

The *mprobe()* function attempts to perform consistency checks on an allocated heap block.

Consistency checks look for inconsistencies within the block header or in the block trailer byte. They may also detect block overruns.

The level of checking provided depends on which version of the allocator you've linked the application with:

- C library — minimal consistency checking.
- Nondebug version of the `malloc` library — a slightly greater level of consistency checking.
- Debug version of the `malloc` library — extensive consistency checking, with tuning available through the use of the *mallopt()* function.

## Returns:

One of the values of the `mcheck_status` enumeration:

`MCHECK_DISABLED`

Consistency checking isn't currently enabled, or consistency information isn't available for this block.

|             |                                                                        |
|-------------|------------------------------------------------------------------------|
| MCHECK_OK   | There are no inconsistencies in this block.                            |
| MCHECK_HEAD | The block header is corrupted.                                         |
| MCHECK_TAIL | The block trailer byte is corrupted or there has been a block overrun. |
| MCHECK_FREE | The <i>ptr</i> argument doesn't point to an allocated heap block.      |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mallopt()*, *mcheck()*

The Heap Analysis chapter in the *Programmer's Guide*.

## Synopsis:

```
#include <sys/mman.h>

int mprotect(const void * addr,
 size_t len,
 int prot);
```

## Arguments:

*addr*     The beginning of the range of addresses whose protection you want to change.

*len*     The length of the range of addresses, in bytes.

*prot*     The new access capabilities for the mapped memory region(s). You can combine the following bits, which are defined in `<sys/mman.h>`:

- PROT\_EXEC — the region can be executed.
- PROT\_NOCACHE — disable caching of the region (for example, to access dual ported memory).
- PROT\_NONE — the region can't be accessed.
- PROT\_READ — the region can be read.
- PROT\_WRITE — the region can be written.

## Library:

`libc`

## Description:

The `mprotect()` function changes the access protections on any mappings residing in the range starting at *addr*, and continuing for *len* bytes.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).



---

If *mprotect()* fails, the protections on some of the pages in the address range starting at *addr* and continuing for *len* bytes may have been changed.

---

**Errors:**

- EACCES The memory object wasn't opened for read, regardless of the protection specified.  
The memory object wasn't opened for write, and PROT\_WRITE was specified for a MAP\_SHARED type mapping.
- EAGAIN The *prot* argument specifies PROT\_WRITE on a MAP\_PRIVATE mapping, and there's insufficient memory resources to reserve for locking the private pages (if required).
- ENOMEM The addresses in the range starting at *addr* and continuing for *len* bytes are outside the range allowed for the address space of a process, or specify one or more pages that are not mapped.  
The *prot* argument specifies PROT\_WRITE on a MAP\_PRIVATE mapping, and locking the private pages (if required) would need more space than the system can supply to reserve for doing so.
- ENOSYS The function *mprotect()* isn't supported by this implementation.



## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*mmap()*, *munmap()*, *shm\_open()*, *shm\_unlink()*

## ***mq\_close()***

© 2004, QNX Software Systems Ltd.

*Close a message queue*

### **Synopsis:**

```
#include <mqueue.h>

int mq_close(mqd_t mqdes);
```

### **Arguments:**

*mqdes*     The message-queue descriptor, returned by *mq\_open()*, of the message queue that you want to close.

### **Library:**

`libc`

### **Description:**

The *mq\_close()* function removes the association between *mqdes* and a message queue. If the current process attaches a notify to this queue for notification, the attachment is eliminated. If this queue is unlinked before the call to *mq\_close()*, and this process is the last process to call *mq\_close()* on the queue, then the queue is destroyed, along with its contents.

Calling *close()* with *mqdes* has the same effect as calling *mq\_close()*.

### **Returns:**

-1 if an error occurred (*errno* is set). Any other value indicates success.

### **Errors:**

EBADF     Invalid queue *mqdes*.

### **Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mq\_open(), mq\_unlink()*

## ***mq\_getattr()***

© 2004, QNX Software Systems Ltd.

*Get a message queue's attributes*

### **Synopsis:**

```
#include <mqueue.h>

int mq_getattr(mqd_t mqdes,
 struct mq_attr* mqstat);
```

### **Arguments:**

*mqdes*      The message-queue descriptor, returned by *mq\_open()*, of the message queue that you want to get the attributes of.

*mqstat*     A pointer to a **mq\_attr** structure where the function can store the attributes of the message queue. For more information, see below.

### **Library:**

**libc**

### **Description:**

The *mq\_getattr()* function determines the current attributes of the queue referenced by *mqdes*. These attributes are stored in the location pointed to by *mqstat*. The fields of the **mq\_attr** structure are as follows:

**long mq\_flags**      The options set for this open message-queue description (i.e. these options are for the given *mqdes*, not the queue as a whole). This field may have been changed by call to *mq\_setattr()* since you opened the queue.

- **O\_NONBLOCK** — no call to *mq\_receive()* or *mq\_send()* will ever block on this queue. If the queue is in such a condition that the given operation can't be performed without blocking, then an error is returned, and *errno* is set to **EAGAIN**.

This constant is specified by POSIX 1003.1b.

**long** *mq\_maxmsg*

The maximum number of messages that can be stored on the queue. This value was set when the queue was created.

**long** *mq\_msgsize*

The maximum size of each message on the given message queue. This value was also set when the queue was created.

**long** *mq\_curmsgs*

The number of messages currently on the given queue.

**long** *mq\_sendwait*

The number of threads currently waiting to send a message. This field was eliminated from the POSIX standard after draft 9, but has been resurrected for the QNX implementation. A nonzero value in this field implies that the queue is full.

**long** *mq\_rcvwait*

The number of threads currently waiting to receive a message. Like *mq\_sendwait*, this field was resurrected for the QNX implementation. A nonzero value in this field implies that the queue is empty.

**Returns:**

-1 if an error occurred (*errno* is set). Any other value indicates success.

**Errors:**

EBADF      Invalid message queue *mqdes*.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mq\_close()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*

### Synopsis:

```
#include <mqueue.h>

int mq_notify(
 mqd_t mqdes,
 const struct sigevent* notification);
```

### Arguments:

|                     |                                                                                                                          |
|---------------------|--------------------------------------------------------------------------------------------------------------------------|
| <i>mqdes</i>        | The message-queue descriptor, returned by <i>mq_open()</i> , of the message queue that you want to get notification for. |
| <i>notification</i> | NULL, or a pointer to a <b>sigevent</b> structure that describes how you want to be notified.                            |

### Library:

**libc**

### Description:

If *notification* isn't NULL, the *mq\_notify()* function asks the server to notify the calling process when the queue makes the transition from empty to nonempty. The means by which the server is to notify the process is passed in the **sigevent** structure pointed to by *notification*. Once the message queue server has notified the process of the transition, the notification is removed.

We recommend that you use the following event types in this case:

- SIGEV\_SIGNAL
- SIGEV\_SIGNAL\_CODE
- SIGEV\_SIGNAL\_THREAD
- SIGEV\_PULSE
- SIGEV\_INTR

Under normal operation, only one process may register for notification at a time. If a process attempts to attach a notification, and another process is already attached, an error is returned and *errno* is set to EBUSY.

If a process has registered for notification, and another process is blocked on *mq\_receive()*, then the *mq\_receive()* call is satisfied by any arriving message. The resulting behavior is as if the message queue remained empty.

If *notification* is NULL and the current process is currently registered for notification, then the existing registration is removed.

**Returns:**

-1 if an error occurred (*errno* is set). Any other value indicates success.

**Errors:**

- EBADF      Invalid message queue *mqdes*.
- EBUSY      A process has already registered for notification for the given queue.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



**See also:**

*mq\_open()*, *mq\_receive()*, *mq\_send()*, **sigevent**

## ***mq\_open()***

*Open a message queue*

© 2004, QNX Software Systems Ltd.

### **Synopsis:**

```
#include <mqueue.h>

mqd_t mq_open(const char * name,
 int oflag,
 ...)
```

### **Arguments:**

- name*      The name of the message queue that you want to open; see below.
- oflag*      You must specify one of O\_RDONLY (receive-only), O\_WRONLY (send-only) or O\_RDWR (send-receive). In addition, you can OR in the following constants to produce the following effects:
- O\_CREAT — if *name* doesn't exist, instruct the server to create a new message queue with the given name. If you specify this flag, *mq\_open()* uses its *mode* and *mq\_attr* arguments; see below.
  - O\_EXCL — if you set both O\_EXCL and O\_CREAT, and a message queue *name* exists, the call fails and *errno* is set to EEXIST. Otherwise, the queue is created normally. If you set O\_EXCL without O\_CREAT, it's ignored.
  - O\_NONBLOCK — under normal message queue operation, a call to *mq\_send()* or *mq\_receive()* could block if the message queue is full or empty. If you set this flag, these calls never block. If the queue isn't in a condition to perform the given call, *errno* is set to EAGAIN and the call returns an error.

If you set O\_CREAT in the *oflag* argument, you must also pass these arguments to *mq\_open()*:

- mode*      The file permissions for the new queue. For more information, see "Access permissions" in the documentation for *stat()*.

If you set any bits other than file permission bits, they're ignored. Read and write permissions are analogous to receive and send permissions; execute permissions are ignored.

*mq\_attr* NULL, or a pointer to an **mq\_attr** structure that contains the attributes that you want to use for the new queue. For more information, see *mq\_getattr()*.

If *mq\_attr* is NULL, the following default attributes are used (provided that no defaults were specified when starting the message queue server):

- *mq\_maxmsg*: 1024
- *mq\_msgsize*: 4096
- *mq\_flags*: 0

If *mq\_attr* isn't NULL, the new queue adopts the *mq\_maxmsg* and *mq\_msgsize* of **mq\_attr**. The *mq\_flags* flags field is ignored.

## Library:

**libc**

## Description:

The *mq\_open()* function opens a message queue referred to by *name*, and returns a message queue descriptor by which the queue can be referenced in the future. The name is interpreted as follows:

| <i>name</i>            | Pathname space entry      |
|------------------------|---------------------------|
| <i>entry</i>           | <i>CWD/entry</i>          |
| <i>/entry</i>          | <i>/dev/mqueue/entry</i>  |
| <i>entry/newentry</i>  | <i>CWD/entry/newentry</i> |
| <i>/entry/newentry</i> | <i>/entry/newentry</i>    |

where *CWD* is the current working directory for the program at the point that it calls *mq\_open()*.



---

If you want to open a queue on another node, you have to specify the name as `/net/node/mqueue_location`.

---

If *name* doesn't exist, *mq\_open()* examines the third and fourth parameters: a `mode_t` and a pointer to an `mq_attr` structure.

The only time that a call to *mq\_open()* with `O_CREAT` set fails is if you open a message queue and later unlink it, but never close it. Like their file counterparts, an unlinked queue that hasn't yet been closed must continue to exist; an attempt to recreate such a message queue fails, and *errno* is set to `ENOENT`.



---

Message queues persist — like files — even after the processes that created them end. A message queue is destroyed when the last process connected to it unlinks from the queue by calling *mq\_unlink()*.

---

## Returns:

A valid message queue descriptor if the queue is successfully created, or -1 (*errno* is set).

## Errors:

|        |                                                                                                                                                                                           |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES | The message queue exists, and you don't have permission to open the queue under the given <i>oflag</i> , or the message queue doesn't exist, and you don't have permission to create one. |
| EEXIST | You specified the <code>O_CREAT</code> and <code>O_EXCL</code> flags in <i>oflag</i> , and the queue <i>name</i> exists.                                                                  |
| EINTR  | The operation was interrupted by a signal.                                                                                                                                                |

|              |                                                                                                                                                |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| EINVAL       | You specified the O_CREAT flag in <i>oflag</i> , and <i>mq_attr</i> wasn't NULL, but some values in the <b>mq_attr</b> structure were invalid. |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                 |
| EMFILE       | Too many file descriptors are in use by the calling process.                                                                                   |
| ENAMETOOLONG | The length of <i>name</i> exceeds PATH_MAX.                                                                                                    |
| ENOENT       | You didn't set the O_CREAT flag, and the queue <i>name</i> doesn't exist.                                                                      |
| ENOSPC       | The message queue server has run out of memory.                                                                                                |
| ENOSYS       | The <i>mq_open()</i> function isn't implemented for the filesystem specified in <i>name</i> .                                                  |

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*mq\_close()*, *mq\_getattr()*, *mq\_notify()*, *mq\_receive()*, *mq\_send()*, *mq\_setattr()*, *mq\_timedreceive()*, *mq\_timedsend()*, *mq\_unlink()*

## ***mq\_receive()***

© 2004, QNX Software Systems Ltd.

*Receive a message from a queue*

### **Synopsis:**

```
#include <mqueue.h>

ssize_t mq_receive(mqd_t mqdes,
 char* msg_ptr,
 size_t msg_len,
 unsigned int* msg_prio);
```

### **Arguments:**

|                 |                                                                                                                        |
|-----------------|------------------------------------------------------------------------------------------------------------------------|
| <i>mqdes</i>    | The message-queue descriptor, returned by <i>mq_open()</i> , of the message queue that you want to get a message from. |
| <i>msg_ptr</i>  | A pointer to a buffer where the function can store the message received.                                               |
| <i>msg_len</i>  | The message size of the given queue.                                                                                   |
| <i>msg_prio</i> | NULL, or a pointer to a location where the function can store the priority of the message received.                    |

### **Library:**

libc

### **Description:**

The *mq\_receive()* function is used to receive the oldest of the highest priority messages in the queue specified by *mqdes*. The priority of the message received is put in the location pointed to by *msg\_prio*, the data itself in the location pointed to by *msg\_ptr*, and the size received is be returned.

If you call *mq\_receive()* with a *msg\_len* of anything other than the *mq\_msgsize* of the specified queue, then *mq\_receive()* returns an error, and *errno* is set to EINVAL.

If there are no messages on the queue specified, and O\_NONBLOCK wasn't set in *oflag* during *mq\_open()*, and MQ\_NONBLOCK isn't

present in the queue's *mq\_flags*, then the *mq\_receive()* call blocks. If multiple *mq\_receive()* calls are blocked on a single queue, then they're unblocked in FIFO order as messages arrive.

Calling *read()* with *mqdes* is analogous to calling *mq\_receive()* with a NULL *msg\_prio*.

## Returns:

The size of the message removed from the queue. If the call fails, -1 is returned as the size, no message is removed from the queue, and *errno* is set.

## Errors:

|          |                                                                                                                                                                                                              |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN   | The O_NONBLOCK or MQ_NONBLOCK flags were set and there are no messages currently on the specified queue.                                                                                                     |
| EBADF    | The <i>mqdes</i> argument doesn't represent a valid queue open for reading.                                                                                                                                  |
| EINTR    | The operation was interrupted by a signal.                                                                                                                                                                   |
| EINVAL   | The <i>msg_ptr</i> argument isn't a valid pointer, or <i>msg_len</i> is less than 0, or <i>msg_len</i> is less than the message size specified in <i>mq_open()</i> . The default message size is 4096 bytes. |
| EMSGSIZE | The given <i>msg_len</i> is shorter than the <i>mq_msgsize</i> for the given queue <i>or</i> the given <i>msg_len</i> is too short for the message that would have been received.                            |

## Classification:

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mq\_close()*, *mq\_open()*, *mq\_send()*, *mq\_timedreceive()*, *read()*



**Synopsis:**

```
#include <mqueue.h>

int mq_send(mqd_t mqdes,
 const char * msg_ptr,
 size_t msg_len,
 unsigned int msg_prio);
```

**Arguments:**

|                 |                                                                                                                       |
|-----------------|-----------------------------------------------------------------------------------------------------------------------|
| <i>mqdes</i>    | The message-queue descriptor, returned by <i>mq_open()</i> , of the message queue that you want to send a message to. |
| <i>msg_ptr</i>  | A pointer to the message that you want to send.                                                                       |
| <i>msg_len</i>  | The size of the message.                                                                                              |
| <i>msg_prio</i> | The priority of the message, in the range from 0 to (MQ_PRIO_MAX-1).                                                  |

**Library:**

`libc`

**Description:**

The *mq\_send()* function puts a message of size *msg\_len* and pointed to by *msg\_ptr* into the queue indicated by *mqdes*. The new message has a priority of *msg\_prio*.

The queue is maintained in priority order, and in FIFO order within the same priority.

If the number of elements on the specified queue is equal to its *mq\_maxmsg*, and `O_NONBLOCK` (in *oflag* of *mq\_open()*) has been set, the call to *mq\_send()* blocks. It becomes unblocked when there's room on the queue to send the given message. If more than one *mq\_send()* is blocked on a given queue, and space becomes available

in that queue to send, then the *mq\_send()* with the highest priority message is unblocked.

Calling *write()* with *mqdes* is analogous to calling *mq\_send()* with a *msg\_prio* of 0.

## Returns:

-1 if an error occurred (*errno* is set). Any other value indicates success.

## Errors:

|          |                                                                                                                                                                                                                        |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN   | The O_NONBLOCK flag was set when opening the queue, and the specified queue is full.                                                                                                                                   |
| EBADF    | The <i>mqdes</i> argument doesn't represent a valid message queue descriptor, or <i>mqdes</i> wasn't opened for writing.                                                                                               |
| EINTR    | The call was interrupted by a signal.                                                                                                                                                                                  |
| EINVAL   | One of the following cases is true: <ul style="list-style-type: none"><li>• <i>msg_len</i> was negative</li><li>• <i>msg_prio</i> was greater than (MQ_PRIO_MAX-1)</li><li>• <i>msg_prio</i> was less than 0</li></ul> |
| EMSGSIZE | The <i>msg_len</i> argument was greater than the <i>msgsize</i> associated with the specified queue.                                                                                                                   |

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

Cancellation point Yes

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*mq\_close()*, *mq\_open()*, *mq\_receive()*, *mq\_timedsend()*

## ***mq\_setattr()***

© 2004, QNX Software Systems Ltd.

*Set a queue's attributes*

### **Synopsis:**

```
#include <mqueue.h>

int mq_setattr(mqd_t mqdes,
 const struct mq_attr* mqstat,
 struct mq_attr* omqstat);
```

### **Arguments:**

*mqdes* The message-queue descriptor, returned by *mq\_open()*, of the message queue that you want to set the attributes of.

*mqstat* A pointer to a **mq\_attr** structure that specifies the attributes that you want to use for the message queue. For more information about this structure, see *mq\_getattr()*; for information about which attributes you can set, see below.

*omqstat* NULL, or a pointer to a **mq\_attr** structure where the function can store the old attributes of the message queue.

### **Library:**

libc

### **Description:**

The *mq\_setattr()* function sets the *mq\_flags* field for the specified queue (passed as the *mq\_flags* field in *mqstat*). If *omqstat* isn't NULL, then the old attribute structure is stored in the location that it points to.

This function ignores the *mq\_maxmsg*, *mq\_msgsize*, and *mq\_curmsgs* fields of *mqstat*. The *mq\_flags* field is the bit-wise OR of zero or more of the following constants:

O\_NONBLOCK

No *mq\_receive()* or *mq\_send()* will ever block on this queue. If the queue is in such a condition that the given operation can't be

performed without blocking, then an error is returned, and *errno* is set to EAGAIN.

This constant is specified by POSIX 1003.1b.



---

The settings that you make for *mq\_flags* apply only to the given message-queue description (i.e. locally), not to the queue itself.

---

### Returns:

-1 if the function couldn't change the attributes (*errno* is set). Any other value indicates success.

### Errors:

EBADF      Invalid message queue *mqdes*.

### Classification:

POSIX 1003.1 (Realtime Extensions)

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*mq\_getattr()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*

## ***mq\_timedreceive()***

© 2004, QNX Software Systems Ltd.

*Receive a message from a message queue*

### **Synopsis:**

```
#include <mqueue.h>
#include <time.h>

ssize_t mq_timedreceive(
 mqd_t mqdes,
 char * msg_ptr,
 size_t msg_len,
 unsigned int * msg_prio,
 const struct timespec * abs_timeout);
```

### **Arguments:**

|                    |                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>mqdes</i>       | The descriptor of the message queue you want to receive a message from, returned by <i>mq_open()</i> .                                                                                     |
| <i>msg_ptr</i>     | A pointer to a buffer where the function can store the message data.                                                                                                                       |
| <i>msg_len</i>     | The size of the buffer, in bytes.                                                                                                                                                          |
| <i>msg_prio</i>    | NULL, or a pointer to a location where the function can store the priority of the message that it removed from the queue.                                                                  |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> structure that specifies the <i>absolute</i> time (not the relative time to the current time) to wait before the function stops trying to receive messages. |

### **Library:**

**libc**

### **Description:**

The *mq\_timedreceive()* function receives the oldest of the highest priority messages in the queue specified by *mqdes*.

If you call *mq\_timedreceive()* with a *msg\_len* of anything other than the *mq\_msgsize* of the specified queue, then *mq\_timedreceive()* returns an error, and *errno* is set to EINVAL.

If there are no messages on the queue specified, and O\_NONBLOCK isn't set in *oflag* during *mq\_open()*, and MQ\_NONBLOCK isn't present in the queue's *mq\_flags*, then the *mq\_timedreceive()* call blocks. If multiple *mq\_timedreceive()* calls are blocked on a single queue, then they're unblocked in FIFO order as messages arrive.

Calling *read()* with *mqdes* is analogous to calling *mq\_timedreceive()* with a NULL *msg\_prio*.

## Returns:

The size of the message removed from the queue, or -1 if an error occurred (no message is removed from the queue, and *errno* is set).

## Errors:

|           |                                                                                                                                                                                                              |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN    | The O_NONBLOCK or MQ_NONBLOCK flags were set and there are no messages currently on the specified queue.                                                                                                     |
| EBADF     | The <i>mqdes</i> argument doesn't represent a valid queue open for reading.                                                                                                                                  |
| EINTR     | The operation was interrupted by a signal.                                                                                                                                                                   |
| EINVAL    | The <i>msg_ptr</i> argument isn't a valid pointer, or <i>msg_len</i> is less than 0, or <i>msg_len</i> is less than the message size specified in <i>mq_open()</i> . The default message size is 4096 bytes. |
| EMSGSIZE  | The given <i>msg_len</i> is shorter than the <i>mq_msgsize</i> for the given queue or the given <i>msg_len</i> is too short for the message that would have been received.                                   |
| ETIMEDOUT | The timeout value was exceeded.                                                                                                                                                                              |

## Examples:

Specify an absolute timeout of 1 second:

```
struct timespec tm;

clock_gettime(CLOCK_REALTIME, &tm);
tm.tv_sec += 1;
if(0 > mq_timedreceive(fd, buf, 4096, NULL, t)) {
 ...
}
```

## Classification:

POSIX 1003.1d (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*mq\_close()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_timedsend()*,  
**timespec**



**Synopsis:**

```
#include <mqueue.h>
#include <time.h>

int mq_timedsend(mqd_t mqdes,
 const char * msg_ptr,
 size_t msg_len,
 unsigned int msg_prio,
 const struct timespec * abs_timeout);
```

**Arguments:**

|                    |                                                                                                                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>mqdes</i>       | The descriptor of the message queue you want to put the message into, returned by <i>mq_open()</i> .                                                                                       |
| <i>msg_ptr</i>     | A pointer to the message data.                                                                                                                                                             |
| <i>msg_len</i>     | The size of the buffer, in bytes.                                                                                                                                                          |
| <i>msg_prio</i>    | The priority of the message, in the range from 0 to (MQ_PRIO_MAX-1).                                                                                                                       |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> structure that specifies the <i>absolute</i> time (not the relative time to the current time) to wait before the function stops trying to receive messages. |

**Library:**

**libc**

**Description:**

The *mq\_timedsend()* function puts a message of size *msg\_len* and pointed to by *msg\_ptr* into the queue indicated by *mqdes*. The new message has a priority of *msg\_prio*.

The queue maintained is in priority order, and in FIFO order within the same priority.

If the number of elements on the specified queue is equal to its *mq\_maxmsg*, and neither `O_NONBLOCK` (in *oflag* of *mq\_open()*) nor `MQ_NONBLOCK` (in the queue's *mq\_flags*) has been set, the call to *mq\_timedsend()* blocks. It becomes unblocked when there's room on the queue to send the given message. If more than one *mq\_timedsend()* is blocked on a given queue, and space becomes available in that queue to send, then the *mq\_timedsend()* with the highest priority message is unblocked.

Calling *write()* with *mqdes* is analogous to calling *mq\_timedsend()* with a *msg\_prio* of 0.

## Returns:

-1 if an error occurred (*errno* is set). Any other value indicates success.

## Errors:

|        |                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The <code>O_NONBLOCK</code> flag is set when opening the queue, or the <code>MQ_NONBLOCK</code> flag is set in its attributes, and the specified queue is full.                                                                                                                                                                                                                                                          |
| EBADF  | The <i>mqdes</i> argument doesn't represent a valid message queue descriptor, or <i>mqdes</i> isn't opened for writing.                                                                                                                                                                                                                                                                                                  |
| EINTR  | The call was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                                                                    |
| EINVAL | One of the following is true: <ul style="list-style-type: none"><li>• The <i>msg_len</i> is negative.</li><li>• The <i>msg_prio</i> is greater than <code>(MQ_PRIO_MAX-1)</code>.</li><li>• The <i>msg_prio</i> is less than 0.</li><li>• The <code>MQ_PRIO_RESTRICT</code> flag is set in the <i>mq_attr</i> member of <i>mqdes</i>, and <i>msg_prio</i> is greater than the priority of the calling process.</li></ul> |

|           |                                                                                                     |
|-----------|-----------------------------------------------------------------------------------------------------|
| EMSGSIZE  | The <i>msg_len</i> argument is greater than the <i>msgsize</i> associated with the specified queue. |
| ETIMEDOUT | The timeout value was exceeded.                                                                     |

### Examples:

See the example for *mq\_timedreceive()*.

### Classification:

POSIX 1003.1d (draft)

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*mq\_close()*, *mq\_open()*, *mq\_receive()*, *mq\_send()*, *mq\_timedreceive()*,  
**timespec**

## ***mq\_unlink()***

© 2004, QNX Software Systems Ltd.

*Remove a queue*

---

### **Synopsis:**

```
#include <mqueue.h>

int mq_unlink(const char* name);
```

### **Arguments:**

*name*     The name of the message queue that you want to unlink.

### **Library:**

`libc`

### **Description:**

The *mq\_unlink()* function removes the queue with the given *name*.

If some process has the queue open when the call to *mq\_unlink()* is made, then the actual deletion of the queue is postponed until it has been closed. If a queue exists in the netherworld between unlinking and the actual removal of the queue, then *all* calls to open a queue with the given name fail (even if `O_CREAT` is present in *oflag*). Once the queue is deleted, all elements currently on it are freed. Due to the lazy deletion of queues, it's impossible for any process to be blocked on the message queue when it's deleted.

Calling *unlink()* with a name that resolves to the message queue server's namespace (e.g. `/dev/mqueue/my_queue`) is analogous to calling *mq\_unlink()* with *name* set to the last elements of the pathname (e.g. `my_queue`).

### **Returns:**

-1 if the queue wasn't successfully unlinked (*errno* is set). Any other value indicates that the queue was successfully unlinked.

## Errors:

|              |                                                                                                 |
|--------------|-------------------------------------------------------------------------------------------------|
| EACCES       | You don't have permission to unlink the specified queue.                                        |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                  |
| ENAMETOOLONG | The length of <i>name</i> exceeds PATH_MAX.                                                     |
| ENOENT       | The queue <i>name</i> doesn't exist.                                                            |
| ENOSYS       | The <i>mq_unlink()</i> function isn't implemented for the filesystem specified in <i>path</i> . |

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*mq\_close()*, *mq\_open()*, *unlink()*

## ***mrnd48()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random signed long integer*

### **Synopsis:**

```
#include <stdlib.h>

long mrnd48(void);
```

### **Library:**

libc

### **Description:**

The *mrnd48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a signed **long** integer uniformly distributed over the interval  $[-2^{31}, 2^{31})$ .

Call one of *lcong48()*, *seed48()*, or *srand48()* to initialize the random-number generator before calling *drand48()*, *lrand48()*, or *mrnd48()*.

The *jrand48()* function is a thread-safe version of *mrnd48()*.

### **Returns:**

A pseudo-random **long** integer.

### **Classification:**

Standard Unix

#### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*drand48(), erand48(), jrand48(), lcong48(), lrnd48(), nrand48(), seed48(), srand48()*

## **\_msg\_info**

© 2004, QNX Software Systems Ltd.

*Information about a message*

### **Synopsis:**

```
struct _msg_info { /* _msg_info _server_info */
 uint32_t nd; /* client server */
 uint32_t srcnd; /* server n/a */
 pid_t pid; /* client server */
 int32_t tid; /* thread n/a */
 int32_t chid; /* server server */
 int32_t scoid; /* server server */
 int32_t coid; /* client client */
 int32_t msglen; /* msg n/a */
 int32_t srcmsglen; /* thread n/a */
 int32_t dstmsglen; /* thread n/a */
 int16_t priority; /* thread n/a */
 int16_t flags; /* n/a client */
 uint32_t reserved;
};
```

### **Description:**

The `_msg_info` structure contains information about a message. The members include:

|               |                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------|
| <i>nd</i>     | The node descriptor of the client machine as viewed by the server. See “Node descriptors,” below. |
| <i>srcnd</i>  | The node descriptor of the server, as viewed by the client.                                       |
| <i>pid</i>    | The process ID of the sending thread.                                                             |
| <i>tid</i>    | The thread ID of the sending thread.                                                              |
| <i>chid</i>   | The channel ID that the message was received on.                                                  |
| <i>scoid</i>  | The server connection ID.                                                                         |
| <i>coid</i>   | The client connection ID.                                                                         |
| <i>msglen</i> | The number of bytes received.                                                                     |



|                  |                                                                                                                                                                                                                                                                                                          |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>srcmsglen</i> | The length of the source message, in bytes, as sent by <i>MsgSend*()</i> . This may be greater than the value in <i>msglen</i> . This member is valid only if you set <code>_NTO_CHF_SENDER_LEN</code> in the <i>flags</i> argument to <i>ChannelCreate()</i> for the channel that received the message. |
| <i>dstmsglen</i> | The length of the client's reply buffer, in bytes, as passed to <i>MsgSend*()</i> . This member is valid only if you set <code>_NTO_CHF_REPLY_LEN</code> in the <i>flags</i> argument to <i>ChannelCreate()</i> for the channel that received the message.                                               |
| <i>priority</i>  | The priority of the sending thread.                                                                                                                                                                                                                                                                      |
| <i>flags</i>     | The client has an unblock pending <code>_NTO_ML_UNBLOCK_REQ</code> (i.e. a timeout on the send occurred or a signal was delivered and <code>_NTO_CHF_UNBLOCK</code> is set on the channel).                                                                                                              |




---

The *msglen* and *srcmsglen* members are valid only until the next call to *MsgRead\*()* or *MsgWrite\*()*.

---

If *msglen* is less than *srcmsglen* and is also less than the receive buffer size, the message is a network transaction that requires more reading of data with *MsgRead\*()*.

## Node descriptors

The *nd* (node descriptor) is a temporary numeric description of a remote node. For more information, see the Qnet Networking chapter of the *System Architecture* guide.

| <b>To:</b>                    | <b>Use this function:</b> |
|-------------------------------|---------------------------|
| Compare two <i>nd</i> objects | <i>ND_NODE_CMP()</i>      |

*continued...*

| <b>To:</b>                  | <b>Use this function:</b> |
|-----------------------------|---------------------------|
| Convert a <i>nd</i> to text | <i>netmgr_ndtostr()</i>   |
| Convert text to a <i>nd</i> | <i>netmgr_strtond()</i>   |

**Classification:**

QNX Neutrino

**See also:**

*MsgInfo()*, *MsgReceive()*, *MsgReceivev()*, *ND\_NODE\_CMP()*,  
*netmgr\_ndtostr()*, *netmgr\_remote\_nd()*, *netmgr\_strtond()*

## ***MsgDeliverEvent(), MsgDeliverEvent\_r()***

*Deliver an event through a channel*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgDeliverEvent(int rvid,
 const struct sigevent* event);

int MsgDeliverEvent_r(
 int rvid,
 const struct sigevent* event);
```

### **Arguments:**

*rvid*      The value returned to the server when it receives a message from a client using *MsgReceive\*()*.

*event*     A pointer to a **sigevent** structure that contains the event you want to send. These events are defined in **<sys/siginfo.h>**. The type of event is placed in *event.sigev\_notify*.

### **Library:**

**libc**

### **Description:**

The *MsgDeliverEvent()* and *MsgDeliverEvent\_r()* kernel calls deliver an *event* from a server to a client through a channel connection. They're typically used to perform async IO and async event notification to clients that don't want to block on a server.

These functions are identical except in the way they indicate errors. See the Returns section for details.

Although the server can explicitly send any event it desires, it's more typical for the server to receive a **struct sigevent** in a message from the client that already contains this data. The message also contains information for the server indicating the conditions on when to notify the client with the event. The server then saves the *rvid*


from *MsgReceive\*()* and the event from the message without needing to examine the event in any way. When the trigger conditions are met in the server, such as data becoming available, the server calls *MsgDeliverEvent()* with the saved *rvid* and *event*.

You can use the *SIGEV\_SIGNAL* set of notifications to create an asynchronous design in which the client is interrupted when the event occurs. The client can make this synchronous by using the *SignalWaitinfo()* kernel call to wait for the signal. Where possible, you should use an event-driven synchronous design that's based on *SIGEV\_PULSE*. In this case, the client sends messages to servers, and requests event notification via a pulse.

You're not likely to use the event types *SIGEV\_UNBLOCK* and *SIGEV\_INTR* with this call.

You should use *MsgDeliverEvent()* when two processes need to communicate with each other without the possibility of deadlock. The blocking nature of *MsgSend\*()* introduces a hierarchy of processes in which "sends" flow one way and "replies" the other way.

In the following diagram, processes at the A level can send to processes at the B or C level. Processes at the B level can send to the C level but they should never send to the A level. Likewise, processes at the C level can never send to those at the A or B level. To A, B and C are servers. To B, A is a client and C is a server.




---

*A hierarchy of processes.*

These hierarchies are simple to establish and ensure a clean

deadlock-free design. If these rules are broken then deadlock can occur as shown below:



*A deadlock when sending messages improperly among processes.*

There are common situations which require communication to flow backwards through the hierarchy. For example, A sends to B requesting notification when data is available. B immediately replies to A. At some point in the future, B will have the data A requested and will inform A. B can't send a message to A because this might result in deadlock if A decided to send to B at the same time.

The solution is to have B use a nonblocking *MsgDeliverEvent()* to inform A. A receives this pulse and sends a message to B requesting the data. B then replies with the data. This is the basis for asynchronous IO. Clients send to servers and where necessary, servers use pulses to request clients to resend to them as needed. This is illustrated below:

| Message          | Use                  |
|------------------|----------------------|
| A sends to → B   | Async IO request     |
| B replies to → A | Request acknowledged |

*continued...*

---

| Message              | Use                      |
|----------------------|--------------------------|
| B sends pulse to → A | Requested data available |
| A sends to → B       | Request for the data     |
| B replies to → A     | Reply with data          |

---



In client/server designs, you typically use *MsgDeliverEvent()* in the server, and *MsgSendPulse()* in the client.

---

### Blocking states

None. In the network case, lower priority threads may run.

### Native networking

When you use *MsgDeliverEvent()* to communicate across a network, the return code isn't "reliable". In the local case, *MsgDeliverEvent()* always returns a correct success or failure value. But since *MsgDeliverEvent()* *must be nonblocking*, in the networked case, the return value isn't guaranteed to reflect the actual result on the client's node. This is because *MsgDeliverEvent()* would have to block waiting for the communications between the two **npm-qnets**.

Generally, this isn't a problem, because *MsgDeliverEvent()* is for the benefit of the client anyway — if the client no longer exists, then the client obviously doesn't care that it didn't get the event. The server usually delivers the event and then goes about its business, regardless of the success or failure of the event delivery.

### Returns:

The only difference between these functions is the way they indicate errors:

*MsgDeliverEvent()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

## *MsgDeliverEvent\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

|            |                                                                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN     | The kernel has insufficient resources to enqueue the event.                                                                                        |
| EBADF      | The thread indicated by <i>rvid</i> had its connection detached.                                                                                   |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided.                                                                             |
| ESRCH      | The thread indicated by <i>rvid</i> doesn't exist.                                                                                                 |
| ESRVRFAULT | A fault occurred in the server's address space when it tried to write the pulse message to the server's receive message buffer (SIGEV_PULSE only). |

## Examples:

The following example demonstrates how a client can request a server to notify it with a pulse at a later time (in this case, after the server has slept for two seconds). The server side notifies the client using *MsgDeliverEvent()*.

Here's the header file that's used by `client.c` and `server.c`:

```
struct my_msg
{
 short type;
 struct sigevent event;
};

#define MY_PULSE_CODE _PULSE_CODE_MINAVAIL+5
#define MSG_GIVE_PULSE _IO_MAX+4
#define MY_SERV "my_server_name"
```

Here's the client side that fills in a `struct sigevent` and then receives a pulse:

```
/* client.c */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <errno.h>
#include <sys/neutrino.h>
#include <sys/iomsg.h>

#include "my_hdr.h"

int main(int argc, char **argv)
{
 int chid, coid, srv_coid, rcvid;
 struct my_msg msg;
 struct _pulse pulse;

 /* we need a channel to receive the pulse notification on */
 chid = ChannelCreate(0);

 /* and we need a connection to that channel for the pulse to be
 delivered on */
 coid = ConnectAttach(0, 0, chid, _NTO_SIDE_CHANNEL, 0);

 /* fill in the event structure for a pulse */
 SIGEV_PULSE_INIT(&msg.event, coid, SIGEV_PULSE_PRIO_INHERIT,
 MY_PULSE_CODE, 0);
 msg.type = MSG_GIVE_PULSE;

 /* find the server */
 if ((srv_coid = name_open(MY_SERV, 0)) == -1)
 {
 printf("failed to find server, errno %d\n", errno);
 exit(1);
 }

 /* give the pulse event we initialized above to the server for
 later delivery */
 MsgSend(srv_coid, &msg, sizeof(msg), NULL, 0);

 /* wait for the pulse from the server */
 rcvid = MsgReceivePulse(chid, &pulse, sizeof(pulse), NULL);
 printf("got pulse with code %d, waiting for %d\n", pulse.code,
 MY_PULSE_CODE);

 return 0;
}
```



}

Here's the server side that delivers the pulse defined by the `struct sigevent`:

```

/* server.c */

#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <errno.h>
#include <sys/neutrino.h>
#include <sys/iomsg.h>
#include <sys/iofunc.h>
#include <sys/dispatch.h>

#include "my_hdr.h"

int main(int argc, char **argv)
{
 int rcvid;
 struct my_msg msg;
 name_attach_t *attach;

 /* attach the name the client will use to find us */
 /* our channel will be in the attach structure */
 if ((attach = name_attach(NULL, MY_SERV, 0)) == NULL)
 {
 printf("server:failed to attach name, errno %d\n", errno);
 exit(1);
 }

 /* wait for the message from the client */
 rcvid = MsgReceive(attach->chid, &msg, sizeof(msg), NULL);
 MsgReply(rcvid, 0, NULL, 0);
 if (msg.type == MSG_GIVE_PULSE)
 {
 /* wait until it is time to notify the client */
 sleep(2);

 /* deliver notification to client that client requested */
 MsgDeliverEvent(rcvid, &msg.event);
 printf("server:delivered event\n");
 } else
 {
 printf("server: unexpected message \n");
 }

 return 0;
}

```

}

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

In the case of a pulse event, if the server faults on delivery, the pulse is either lost or an error is returned.

## See also:

*MsgReceive()*, *MsgReceivev()*, *MsgSend()*, *MsgSendPulse()*,  
*MsgSendv()*, **sigevent**, *SignalWaitinfo()*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgError(int rvid,
 int error);

int MsgError_r(int rvid,
 int error);
```

### **Arguments:**

*rvid*     The receive ID that *MsgReceive\*()* returned.

*error*    The error code that you want to set for the client.

### **Library:**

`libc`

### **Description:**

The *MsgError()* and *MsgError\_r()* kernel calls unblock the client's *MsgSend\*()* call and set the client's *errno* to *error*. No data is transferred.

If *error* is EOK, the *MsgSend\*()* call returns EOK; if *error* is any other value, the *MsgSend\*()* call returns -1.

These functions are identical except in the way they indicate errors. See the Returns section for details.



---

An error number of ERESTART causes the sender to immediately call *MsgSend\*()* again. Since send and receive buffers passed to *MsgSend()* may overlap, you shouldn't use ERESTART after a call to *MsgWrite()*.

---

## Blocking states

None. In the network case, lower priority threads may run.

## Native networking

*MsgError()* has increased latency when you use it to communicate across a network — the server is now writing the error code to its local `npm-qnet`, which may need to communicate with the client's `npm-qnet` to actually transfer the error code.

## Returns:

The only difference between these functions is the way they indicate errors:

|                     |                                                                                                                                                                   |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgError()</i>   | If an error occurs, the function returns -1 and sets <i>errno</i> . Any other value returned indicates success.                                                   |
| <i>MsgError_r()</i> | Returns EOK on success. This function does <b>NOT</b> set <i>errno</i> . If an error occurs, the function returns one of the values listed in the Errors section. |

## Errors:

ESRCH     The thread indicated by *rvid* doesn't exist.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ChannelCreate(), MsgRead(), MsgReady(), MsgReceive(),  
MsgReceivev(), MsgSend(), MsgSendv()*

## ***MsgInfo()*, *MsgInfo\_r()***

© 2004, QNX Software Systems Ltd.

*Get additional information about a message*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgInfo(int rvid,
 struct _msg_info* info);

int MsgInfo_r(int rvid,
 struct _msg_info* info);
```

### **Arguments:**

*rvid*     The return value from *MsgReceive\**().

*info*     A pointer to a `_msg_info` structure where the function can store information about the message.

### **Library:**

`libc`

### **Description:**

The *MsgInfo()* and *MsgInfo\_r()* kernel calls get additional information about a received message and store it in the specified `_msg_info` structure.

These functions are identical, except in the way they indicate errors. See the Returns section for details.



---

The *info->msglen* and *info->srcmsglen* members are valid only until the next call to *MsgRead\**() or *MsgWrite\**() .

---

### **Blocking states**

This call doesn't block.

**Returns:**

The only difference between these functions is the way they indicate errors:

- MsgInfo()* If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.
- MsgInfo\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

**Errors:**

- EFAULT A fault occurred when the kernel tried to access the buffers provided.
- ESRCH The thread indicated by *rvid* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ChannelCreate()*, *\_msg\_info*, *MsgRead()*, *MsgReady()*, *MsgReceive()*, *MsgReceivev()*, *MsgSend()*, *MsgSendv()*,

## ***MsgKeyData()*, *MsgKeyData\_r()***

© 2004, QNX Software Systems Ltd.

*Pass data through a common client*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgKeyData(int rvid,
 int op,
 uint32_t key,
 uint32_t * key2,
 const iov_t * msg,
 int parts);

int MsgKeyData_r(int rvid,
 int op,
 uint32_t key,
 uint32_t * key2,
 const iov_t * msg,
 int parts);
```

### **Arguments:**

- rvid*     The return value from *MsgReceive\**().
- op*       The operation to perform; one of:
- `_NTO_KEYDATA_CALCULATE` — calculate a new key.
  - `_NTO_KEYDATA_VERIFY` — verify the key.
- key*      A private value for key (this can be a value returned by the *rand()* function).
- key2*     A pointer to a key. What the function stores in this location depends on the *op* argument:
- `_NTO_KEYDATA_CALCULATE` — the new key.
  - `_NTO_KEYDATA_VERIFY` — zero if no tampering has occurred.
- msg*      A pointer to a portion of the reply data to be keyed.
- parts*    The number of parts in *msg*.



**Library:**

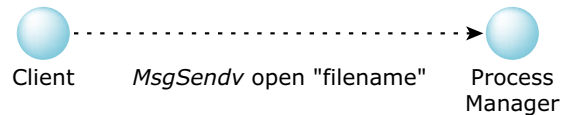
libc

**Description:**

The *MsgKeyData()* and *MsgKeyData\_r()* kernel calls allow two privileged processes to pass data through a common client while verifying that the client hasn't modified the data. This is best explained by an example.

These functions are identical except in the way they indicate errors. See the Returns section for details.

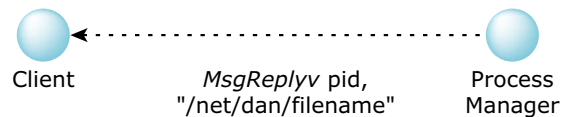
A program calls *open()* with a filename. The *open()* function sends a message to the Process Manager, which is responsible for pathname management.




---

*MsgSendv()*, client to process manager.

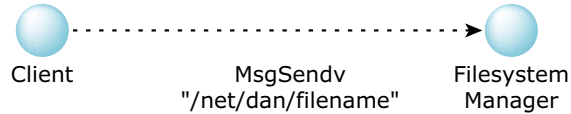
The Process Manager resolves the pathname, resulting in a fully qualified network path and the process ID to send the *open()* request to. This information is replied to the client.




---

*MsgReplyv()*, process manager to client.

The client now sends this message to *pid* with the fully qualified pathname.



*MsgSendv()*, client to filesystem manager

Note that the client can change the pathname before it sends it to the Filesystem Manager. In fact, it could skip the call to the Process Manager and manufacture any pathname it desired. The Filesystem Manager always performs permission checking. Therefore, changing or manufacturing pathnames isn't normally something to be concerned about, except in one case: *chroot()* lets you specify a prefix that must be applied to all pathnames.

In the above example, the client may have had a *chroot()* of `/net/node2/home/dan`. This should limit the process from accessing files outside of `/net/node2/home/dan`. For example:

| User path            | Mapped to <i>chroot()</i> path          |
|----------------------|-----------------------------------------|
| <code>/bin/ls</code> | <code>/net/node2/home/dan/bin/ls</code> |
| <code>/</code>       | <code>/net/node2/home/dan</code>        |

The process has had its root set to a subdirectory, limiting the files it can access. For this to work, it's necessary to prevent the client from changing or manufacturing its own pathnames.



In QNX Neutrino, only the Process Manager handles a user *chroot()*. Unlike a monolithic kernel where the filesystem shares the same address space as the kernel and the *chroot()* information, QNX I/O managers reside in separate address spaces and might not even reside on the same machine.

The solution to this problem is the *MsgKeyData()* call. When the Process Manager receives the *open()* message, it generates the reply data. Before replying, it calls *MsgKeyData()*, with these arguments:

|              |                                                                                                |
|--------------|------------------------------------------------------------------------------------------------|
| <i>rvid</i>  | The return value from <i>MsgReceive*()</i> .                                                   |
| <i>op</i>    | _NTO_KEYDATA_CALCULATE                                                                         |
| <i>key</i>   | A private value for key (this can be a value returned by the <i>rand()</i> function).          |
| <i>key2</i>  | A pointer to a new key that should be returned to the client in a unkeyed area of the message. |
| <i>msg</i>   | A pointer to a portion of the reply data to be keyed.                                          |
| <i>parts</i> | The number of parts in <i>msg</i> .                                                            |

The client now sends the message to the File Manager. On receipt of the message, the File Manager calls *MsgKeyData()* with the same arguments as above, except for:

|            |                                         |
|------------|-----------------------------------------|
| <i>op</i>  | _NTO_KEYDATA_VERIFY                     |
| <i>key</i> | The key that's provided in the message. |

*MsgKeyData()* sets the key pointed to by *key2* to zero if no tampering has occurred.

Note that there are actually two keys involved. A public key that's returned to the client and a private key that the Process Manager generated. The algorithm uses both keys and the data for verification.

### **Blocking states**

These calls don't block.

### **Returns:**

The only difference between these functions is the way they indicate errors:

|                     |                                                                                                         |
|---------------------|---------------------------------------------------------------------------------------------------------|
| <i>MsgKeyData()</i> | If an error occurs, -1 is returned and <i>errno</i> is set. Any other value returned indicates success. |
|---------------------|---------------------------------------------------------------------------------------------------------|

*MsgKeyData\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

ESRCH     The thread indicated by *rvid* doesn't exist.

EFAULT    A fault occurred when the kernel tried to access the buffers provided.

## Examples:

```
/*
 * This program demonstrates the use of MsgKeyData() as a way
 * of a client handing off data from a source server to a
 * destination server such that if the client tampers with
 * the data, the destination server will know about it.
 */

#include <errno.h>
#include <pthread.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <time.h>
#include <unistd.h>
#include <sys/neutrino.h>

typedef struct {
 int public_key;
 char text[10];
} IPC_t;

int chid_src, chid_dst;

void* server_src_thread(void* parm);
void* server_dst_thread(void* parm);

main()
{
 pthread_t tid[2];
 IPC_t msg;
 int coid;
 int status;

 pthread_create(&tid[0], NULL, server_src_thread, NULL);
```

```

pthread_create(&tid[1], NULL, server_dst_thread, NULL);
sleep(3);
/* give time for channels to be created, sloppy but simple */

/*
 * Send to server_src_thread for some data.
 * The data will include some text and a public
 * key for that text.
 */

coid = ConnectAttach(0, 0, chid_src, 0, 0);
MsgSend(coid, NULL, 0, &msg, sizeof(msg));
ConnectDetach(coid);

/*
 * Now send to server_dst_thread with the reply from
 * server_src_thread. We didn't modify the 'text' so it
 * should reply success. Note that we're including the
 * public key.
 */

coid = ConnectAttach(0, 0, chid_dst, 0, 0);
status = MsgSend(coid, &msg, sizeof(msg), &msg, sizeof(msg));
printf("Sent unmodified text to server_dst_thread.
 Replied with %s\n", status == EOK ? "EOK" : "EINVAL");

/*
 * Now tamper with the original 'text' (which we aren't
 * supposed to do) and send to server_dst_thread again
 * but with the modified 'text' and the public key.
 * Since we tampered with the 'text', server_dst_thread
 * should reply failure.
 */

strcpy(msg.text, "NEWDATA");
status = MsgSend(coid, &msg, sizeof(msg), &msg, sizeof(msg));
printf("Sent modified text to server_dst_thread.
 Replied with %s\n", status == EOK ? "EOK" : "EINVAL");

return 0;
}

void* server_src_thread(void* parm)
{
 int rcvid;
 int private_key; /* the kernel keeps this */
 iov_t keyed_area_iov;
 IPC_t msg;
 struct timespec t;

```

```
chid_src = ChannelCreate(0);
while (1) {
 rcvid = MsgReceive(chid_src, &msg, sizeof(msg), NULL);

 /*
 * Give MsgKeyData() the private key and it will
 * calculate a public key for the 'text' member of
 * the message. The kernel will keep the private key
 * and we reply with the public key.
 * Note that we use the number of nanoseconds since the
 * last second as a way of getting a 32-bit pseudo
 * random number for the private key.
 */

 clock_gettime(CLOCK_REALTIME, &t);
 private_key = t.tv_nsec; /* nanoseconds since last second */
 strcpy(msg.text, "OKDATA");
 SETIOV(&keyed_area_iov, &msg.text, sizeof(msg.text));
 MsgKeyData(rcvid, _NTO_KEYDATA_CALCULATE, private_key,
 &msg.public_key, &keyed_area_iov, 1);

 MsgReply(rcvid, 0, &msg, sizeof(msg));
}
return NULL;
}

void* server_dst_thread(void* parm)
{
 int rcvid, tampered, status;
 iov_t keyed_area_iov;
 IPC_t msg;

 chid_dst = ChannelCreate(0);
 while (1) {
 rcvid = MsgReceive(chid_dst, &msg, sizeof(msg), NULL);

 /*
 * Use the public key to see if the data
 * has been tampered with.
 */

 SETIOV(&keyed_area_iov, &msg.text, sizeof(msg.text));
 MsgKeyData(rcvid, _NTO_KEYDATA_VERIFY, msg.public_key,
 &tampered, &keyed_area_iov, 1);

 if (tampered)
 status = EINVAL; /* reply: 'text' was modified */
 else
 status = EOK; /* reply: 'text' was okay */
 MsgReply(rcvid, status, &msg, sizeof(msg));
 }
}
```

```
 }
 return NULL;
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*chroot(), MsgReceive(), MsgReceivev(), open(), rand()*

# ***MsgRead(), MsgRead\_r()***

© 2004, QNX Software Systems Ltd.

*Read data from a message*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgRead(int rvid,
 void* msg,
 int bytes,
 int offset);

int MsgRead_r(int rvid,
 void* msg,
 int bytes,
 int offset);
```

## **Arguments:**

*rvid*      The value returned by *MsgReceive\*()* when you received the message.

*msg*        A pointer to a buffer where the function can store the data.

*bytes*      The number of bytes that you want to read. These functions don't let you read past the end of the thread's message; they return the number of bytes actually read.

*offset*     An offset into the thread's send message that indicates where you want to start reading the data.

## **Library:**

`libc`

## **Description:**

The *MsgRead()* and *MsgRead\_r()* kernel calls read data from a message sent by a thread identified by *rvid*. The thread being read from must not have been replied to and will be in the REPLY-blocked state. Any thread in the receiving process is free to read the message.

These functions are identical except in the way they indicate errors. See the Returns section for details.



The data transfer occurs immediately and the thread doesn't block. The state of the sending thread doesn't change.

You'll use these functions in these situations:

- A message is sent consisting of a fixed header and a variable amount of data. The header contains the byte count of the data. If the data is large and has to be inserted into one or more buffers (like a filesystem cache), rather than read the data into one large buffer and then copy it into several other buffers, *MsgReceive()* reads only the header, and you can call *MsgRead()* one or more times to read data directly into the required buffer(s).
- A message is received but can't be handled at the present time. At some point in the future, an event will occur that will allow the message to be processed. Rather than saving the message until it can be processed (thus using memory resources), you can use *MsgRead()* to reread the message, during which time the sending thread is still blocked.
- Messages that are larger than available buffer space are received. Perhaps the process is an agent between two processes and simply filters the data and passes it on. You can use *MsgRead()* to read the message in small pieces, and use *MsgWrite\*()* to write the messages in small pieces.

When you're finished using *MsgRead()*, you must use *MsgReply\*()* to ready the REPLY-blocked process and complete the message exchange.

### **Blocking states**

None. In the network case, lower priority threads may run.

### **Native networking**

The *MsgRead()* function has increased latency when it's used to communicate across a network — a message pass is involved from the server to the network manager (at least). Depending on the size of the data transfer, the server's **npm-qnet** and the client's **npm-qnet** may

need to communicate over the link to read more data bytes from the client.

## Returns:

The only difference between the *MsgRead()* and *MsgRead\_r()* functions is the way they indicate errors:

*MsgRead()*      The number of bytes read. If an error occurs, -1 is returned and *errno* is set.

*MsgRead\_r()*    The number of bytes read. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

If you try to read past the end of the thread's message, the functions return the number of bytes they were actually able to read.

## Errors:

|            |                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in a server's address space when it tried to access the caller's message buffers. |
| ESRCH      | The thread indicated by <i>rcvid</i> doesn't exist or has had its connection detached.             |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                             |

## Classification:

QNX Neutrino

### Safety

Cancellation point    No

Interrupt handler      No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*MsgReady(), MsgReceive(), MsgReceivev(), MsgReply(),  
MsgReplyv(), MsgWrite(), MsgWritev()*

# ***MsgReadv()*, *MsgReadv\_r()***

© 2004, QNX Software Systems Ltd.

*Read data from a message*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReadv(int rvid,
 const iov_t* riov,
 int rparts,
 int offset);

int MsgReadv_r(int rvid,
 const iov_t* riov,
 int rparts,
 int offset);
```

## **Arguments:**

|               |                                                                                                   |
|---------------|---------------------------------------------------------------------------------------------------|
| <i>rvid</i>   | The value returned by <i>MsgReceive*()</i> when you received the message.                         |
| <i>riov</i>   | An array of buffers where the functions can store the data.                                       |
| <i>rparts</i> | The number of elements in the <i>riov</i> array.                                                  |
| <i>offset</i> | An offset into the thread's send message that indicates where you want to start reading the data. |

## **Library:**

`libc`

## **Description:**

The *MsgReadv()* and *MsgReadv\_r()* kernel calls read data from a message sent by a thread identified by *rvid*. The thread being read from must not have been replied to and will be in the REPLY-blocked state. Any thread in the receiving process is free to read the message.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The data transfer occurs immediately and the thread doesn't block. The state of the sending thread doesn't change.

An attempt to read past the end of the thread's message results in fewer bytes returned than requested.

You'll use these functions in these situations:

- A message is sent consisting of a fixed header and a variable amount of data. The header contains the byte count of the data. If the data is large and has to be inserted into one or more buffers (like a filesystem cache), rather than read the data into one large buffer and then copy it into several other buffers, *MsgReceive()* reads only the header, and you can build a custom `iovec_t` list to let *MsgReadv()* read data directly into the required buffers.
- A message is received but can't be handled at the present time. At some point in the future, an event will occur that will allow the message to be processed. Rather than saving the message until it can be processed (thus using memory resources), you can use *MsgReadv()* to reread the message, during which time the sending thread is still blocked.
- Messages that are larger than available buffer space are received. Perhaps the process is an agent between two processes and simply filters the data and passes it on. You can use *MsgReadv()* to read the message in small pieces, and use *MsgWrite\*()* to write the messages in small pieces.

When you're finished using *MsgReadv()*, you must use *MsgReply\*()* to ready the REPLY-blocked process and complete the message exchange.

### Blocking states

None. In the network case, lower priority threads may run.

### Returns:

The only difference between the *MsgReadv()* and *MsgReadv\_r()* functions is the way they indicate errors:

*MsgReadv()*      The number of bytes read. If an error occurs, -1 is returned and *errno* is set.

*MsgReadv\_r()* The number of bytes read. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

|            |                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in a server's address space when it tried to access the caller's message buffers. |
| ESRCH      | The thread indicated by <i>rcvid</i> doesn't exist or has had its connection detached.             |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                             |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*MsgRead(), MsgReceive(), MsgReceivev(), MsgReply(), MsgReplyv(), MsgWrite(), MsgWritev()*

## ***MsgReceive()*, *MsgReceive\_r()***

*Wait for a message or pulse on a channel*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReceive(int chid,
 void * msg,
 int bytes,
 struct _msg_info * info);

int MsgReceive_r(int chid,
 void * msg,
 int bytes,
 struct _msg_info * info);
```

### **Arguments:**

*chid*      The ID of a channel that you established by calling *ChannelCreate()*.

*msg*        A pointer to a buffer where the function can store the received data.

*bytes*      The size of the buffer.

*info*        NULL, or a pointer to a `_msg_info` structure where the function can store additional information about the message.

### **Library:**

`libc`

### **Description:**

The *MsgReceive()* and *MsgReceive\_r()* kernel calls wait for a message or pulse to arrive on the channel identified by *chid*, and store the received data in the buffer pointed to by *msg*.

These functions are identical, except in the way they indicate errors; see the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The received data isn't allowed to overflow the receive buffer area provided.



---

The *msg* buffer *must* be big enough to contain a pulse. If it isn't, the functions indicate an error of EFAULT.

---

If a message is waiting on the channel when you call *MsgReceive()*, the calling thread doesn't block, and the message is immediately copied. If a message isn't waiting, the calling thread enters the RECEIVE-blocked state until a message arrives.

If multiple messages are sent to a channel without a thread waiting to receive them, the messages are queued in priority order.

If you pass a non-NULL pointer for *info*, the functions store additional information about the message and the thread that sent it in the `_msg_info` structure that *info* points to. You can get this information later by calling *MsgInfo()*.

On success, *MsgReceive()* and *MsgReceive\_r()* return:

- >0 A message was received; the returned value is a *rvid* (receive identifier). You'll use the *rvid* with other *Msg\*()* kernel calls to interact with and reply to the sending thread. *MsgReceive()* changes the state of the sending thread to REPLY-blocked when the message is received. When you use *MsgReply\*()* to reply to the received message, the sending thread is made ready again. The *rvid* encodes the sending thread's ID and a local connection ID.
- 0 A pulse was received; *msg* contains a pulse message of type `_pulse`. When a pulse is received, the kernel space allocated to hold it is immediately released. The `_msg_info` structure isn't updated.





---

Don't reply to a pulse.

---

### Blocking states

| State         | Meaning                    |
|---------------|----------------------------|
| STATE_RECEIVE | There's no message waiting |

### Native networking

In networked message-passing transactions, the most noticeable impact is on the server. The server receives the client's message from the server's local **npm-qnet**. Note that the receive ID that comes back from *MsgReceive()* will have some differences, but you don't need to worry about the format of the receive ID — just treat it as a “magic cookie.”

When the server unblocks from its *MsgReceive()*, it may or may not have received as much of the message as it would in the local case. This is because of the way that message passing is defined — the client and the server agree on the size of the message transfer area (the transmit parameters passed to *MsgSend()* on the client end) and the size of the message receive area on the server's *MsgReceive()*.

In a local message pass, the kernel would ordinarily limit the size of the transfer to the minimum of both sizes. But in the networked case, the message is received by the client's **npm-qnet** into its own private buffers and then sent via transport to the remote **npm-qnet**. Since the size of the server's receive data area can't be known in advance by the client's **npm-qnet** when the message is sent, only a fixed maximum size (currently 8K) message is transferred between the client and the server.

This means, for example, that if the client sends 1 Mbyte of data and the server issues a *MsgReceive()* with a 1-Mbyte data area, then only the number of bytes determined by a network manager would in fact be transferred. The number of bytes transferred to the server is returned via the last parameter to *MsgReceive()* or a call to *MsgInfo()*,

specifically the *msglen* member of struct `_msg_info`. The client doesn't notice this, because it's still blocked.

You can use the following code to ensure that the desired number of bytes are received. Note that this is handled for you automatically when you're using the resource manager library:

```
chid = ChannelCreate(_NTO_CHF_SENDER_LEN);
...
rcvid = MsgReceive(chid, msg, nbytes, &info);

/*
 Doing a network transaction and not all
 the message was send, so get the rest...
*/
if (rcvid > 0 && info.srcmsglen > info.msglen && info.msglen < nbytes) {
 int n;

 if((n = MsgRead_r(rcvid, (char *) msg + info.msglen,
 nbytes - info.msglen, info.msglen)) < 0) {
 MsgError(rcvid, -n);
 continue;
 }
 info.msglen += n;
}
```

## Returns:

The only difference between *MsgReceive()* and *MsgReceive\_r()* is the way they indicate errors. On success, both functions return a positive *rcvid* if they received a message, or 0 if they received a pulse.

If an error occurs:

- *MsgReceive()* returns -1 and sets *errno*.
- *MsgReceive\_r()* returns the negative of a value from the Errors section is returned. This function **doesn't** set *errno*.

**Errors:**

|           |                                                                                                                                                                                                                                                                                                                                                                        |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. Because the OS accesses the sender's buffers only when <i>MsgReceive()</i> is called, a fault could occur <i>in the sender</i> if the sender's buffers are invalid. If a fault occurs when accessing the sender buffers (only) they'll receive an EFAULT and <i>MsgReceive()</i> won't unblock. |
| EINTR     | The call was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                  |
| ESRCH     | The channel indicated by <i>chid</i> doesn't exist.                                                                                                                                                                                                                                                                                                                    |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                                                                                                                                       |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ChannelCreate()*, **`_msg_info`**, *MsgInfo()*, *MsgRead()*, *MsgReadv()*, *MsgReceivePulse()*, *MsgReceivePulsev()*, *MsgReceivev()*, *MsgReply()*, *MsgReplyv()*, *MsgSend()*, *MsgWrite()*, *MsgWritev()*, **`_pulse`**, *TimerTimeout()*

# ***MsgReceivePulse()*, *MsgReceivePulse\_r()***

© 2004, QNX

Software Systems Ltd.

*Receive a pulse on a channel*

---

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReceivePulse(int chid,
 void * pulse,
 int bytes,
 struct _msg_info * info);

int MsgReceivePulse_r(int chid,
 void * pulse,
 int bytes,
 struct _msg_info * info);
```

## **Arguments:**

*chid*      The ID of a channel that you established by calling *ChannelCreate()*.

*pulse*     A pointer to a buffer where the function can store the received data.

*bytes*     The size of the buffer.

*info*      The function doesn't update this structure, so you typically pass NULL for this argument.

## **Library:**

`libc`

## **Description:**

The *MsgReceivePulse()* and *MsgReceivePulse\_r()* kernel calls wait for a pulse to arrive on the channel identified by *chid* and place the received data in the buffer pointed to by *pulse*.

These functions are identical, except in the way they indicate errors; see the Returns section for details.

## *MsgReceivePulse()*, *MsgReceivePulse\_r()*

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The received data isn't allowed to overflow the receive buffer area provided.

---

☞ The *pulse* buffer *must* be big enough to contain a pulse. If it isn't, the functions indicate an error of EFAULT.

---

If a pulse is waiting on the channel when you call *MsgReceivePulse()*, the calling thread doesn't block, and the pulse is immediately copied. If a pulse isn't waiting, the calling thread enters the RECEIVE-blocked state until a pulse arrives.

If multiple pulses are sent to a channel without a thread waiting to receive them, the pulses are queued in priority order.

On success, *MsgReceivePulse()* and *MsgReceivePulse\_r()* return 0 to indicate that they received a pulse. When a pulse is received:

- the kernel space allocated to hold it is immediately released
- *pulse* contains a pulse message of type `_pulse`.

---

☞ Don't reply to a pulse.

---

### Blocking states

| State         | Meaning                   |
|---------------|---------------------------|
| STATE_RECEIVE | There's no pulse waiting. |

### Returns:

The only difference between *MsgReceivePulse()* and *MsgReceivePulse\_r()* is the way they indicate errors. On success, they both return 0.

If an error occurred:

- *MsgReceivePulse()* returns -1 and sets *errno*.

# *MsgReceivePulse()*, *MsgReceivePulse\_r()*

© 2004, QNX

Software Systems Ltd.

---

- *MsgReceivePulse\_r()* returns the negative of a value from the Errors section. This function **doesn't** set *errno*.

## Errors:

|           |                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. Because the OS accesses the sender's buffers only when <i>MsgReceivePulse()</i> is called, a fault could occur <i>in the sender</i> if the sender's buffers are invalid. If a fault occurs when accessing the sender buffers (only) they'll receive an EFAULT and <i>MsgReceivePulse()</i> won't unblock. |
| EINTR     | The call was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                            |
| ESRCH     | The channel indicated by <i>chid</i> doesn't exist.                                                                                                                                                                                                                                                                                                                              |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                                                                                                                                                 |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*MsgDeliverEvent()*, *MsgReceive()*, *MsgReceivePulsev()*,  
*MsgReceivev()*, *MsgSendPulse()*, **\_pulse**, *TimerTimeout()*

## ***MsgReceivePulsev(), MsgReceivePulsev\_r()***

*Receive a pulse on a channel*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReceivePulsev(int chid,
 const iov_t * piov,
 int parts,
 struct _msg_info * info);

int MsgReceivePulsev_r(int chid,
 const iov_t * piov,
 int parts,
 struct _msg_info * info);
```

### **Arguments:**

*chid*      The ID of a channel that you established by calling *ChannelCreate()*.

*piov*      An array of buffers where the function can store the received data.

*parts*     The number of elements in the array.

*info*      The function doesn't update this structure, so you typically pass NULL for this argument.

### **Library:**

`libc`

### **Description:**

The *MsgReceivePulsev()* and *MsgReceivePulsev\_r()* kernel calls wait for a pulse to arrive on the channel identified by *chid* and places the received data in the array of buffers pointed to by *piov*.

These functions are identical, except in the way they indicate errors; see the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The received data isn't allowed to overflow the receive buffer area provided.



---

The first buffer of the IOV (input/output vector) *must* be big enough to contain a pulse. If it isn't, the functions indicate an error of EFAULT.

---

If a pulse is waiting on the channel when you call *MsgReceivePulsev()*, the calling thread doesn't block, and the pulse is immediately copied. If a pulse isn't waiting, the calling thread enters the RECEIVE-blocked state until a pulse arrives.

If multiple pulses are sent to a channel without a thread waiting to receive them, the pulses are queued in priority order.

On success, *MsgReceivePulsev()* and *MsgReceivePulsev\_r()* return 0 to indicate that they received a pulse. When a pulse is received:

- the kernel space allocated to hold it is immediately released
- the IOV's first buffer contains a pulse message of type `_pulse`.

## Blocking states

| State         | Meaning                   |
|---------------|---------------------------|
| STATE_RECEIVE | There's no pulse waiting. |

## Returns:

The only difference between *MsgReceivePulsev()* and *MsgReceivePulsev\_r()* is the way they indicate errors. On success, they both return 0.

If an error occurs:

- *MsgReceivePulsev()* returns -1 and sets *errno*.
- *MsgReceivePulsev\_r()* returns the negative of a value from the Errors section. This function **doesn't** set *errno*.



## Errors:

|           |                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. Because the OS accesses the sender's buffers only when <i>MsgReceivePulsev()</i> is called, a fault could occur <i>in the sender</i> if the sender's buffers are invalid. If a fault occurs when accessing the sender buffers (only) they'll receive an EFAULT and <i>MsgReceivePulsev()</i> won't unblock. |
| EINTR     | The call was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                              |
| ESRCH     | The channel indicated by <i>chid</i> doesn't exist.                                                                                                                                                                                                                                                                                                                                |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                                                                                                                                                   |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*MsgDeliverEvent()*, *MsgReceive()*, *MsgReceivePulse()*,  
*MsgReceivev()*, *MsgSendPulse()*, **`_pulse`**, *TimerTimeout()*

## ***MsgReceivev()*, *MsgReceivev\_r()***

© 2004, QNX Software Systems Ltd.

*Wait for a message or pulse on a channel*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReceivev(int chid,
 const iov_t * riov,
 int rparts,
 struct _msg_info * info);

int MsgReceivev_r(int chid,
 const iov_t * riov,
 int rparts,
 struct _msg_info * info);
```

### **Arguments:**

*chid*      The ID of a channel that you established by calling *ChannelCreate()*.

*riov*      An array of buffers where the function can store the received data.

*rparts*    The number of elements in the array.

*info*      NULL, or a pointer to a `_msg_info` structure where the function can store additional information about the message.

### **Library:**

`libc`

### **Description:**

The *MsgReceivev()* and *MsgReceivev\_r()* kernel calls wait for a message or pulse to arrive on the channel identified by *chid* and place the received data in the array of buffers pointed to by *riov*.

These functions are identical, except in the way they indicate errors; see the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The received data isn't allowed to overflow the receive buffer area provided.



---

The first buffer of the IOV (input/output vector) *must* be big enough to contain a pulse. If it isn't, the functions indicate an error of EFAULT.

---

If a message is waiting on the channel when you call *MsgReceivev()*, the calling thread won't block, and the message is immediately copied. If a message isn't waiting, the calling thread enters the RECEIVE-blocked state until a message arrives.

If multiple messages are sent to a channel without a thread waiting to receive them, the messages are queued in priority order.

If you pass a non-NULL pointer for *info*, the functions store additional information about the message and the thread that sent it in the `_msg_info` structure that *info* points to. You can get this information later by calling *MsgInfo()*.

On success, *MsgReceivev()* and *MsgReceivev\_r()* return:

- >0 A message was received; the value returned is a *rvid* (receive identifier). You'll use the *rvid* with other *Msg\*()* kernel calls to interact with and reply to the sending thread. *MsgReceivev()* changes the state of the sending thread to REPLY-blocked when the message is received. When you use *MsgReply\*()* to reply to the received message, the sending thread is made ready again. The *rvid* encodes the sending thread's ID and a local connection ID.
- 0 A pulse was received; the IOV's first buffer contains a pulse message of type `_pulse`. When a pulse is received, the kernel space allocated to hold it is immediately released. The `_msg_info` structure isn't updated.



---

Don't reply to a pulse.

---

## Blocking states

| State         | Meaning                     |
|---------------|-----------------------------|
| STATE_RECEIVE | There's no message waiting. |

## Returns:

The only difference between *MsgReceivev()* and *MsgReceivev\_r()* is the way they indicate errors. On success, both functions return a positive *rvid* if they received a message, or 0 if they received a pulse.

If an error occurs:

- *MsgReceivev()* returns -1 and sets *errno*.
- *MsgReceivev\_r()* returns the negative of a value from the Errors section. This function **doesn't** set *errno*.

## Errors:

|           |                                                                                                                                                                                                                                                                                                                                                                          |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. Because the OS accesses the sender's buffers only when <i>MsgReceivev()</i> is called, a fault could occur <i>in the sender</i> if the sender's buffers are invalid. If a fault occurs when accessing the sender buffers (only) they'll receive an EFAULT and <i>MsgReceivev()</i> won't unblock. |
| EINTR     | The call was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                    |
| ESRCH     | The channel indicated by <i>chid</i> doesn't exist.                                                                                                                                                                                                                                                                                                                      |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                                                                                                                                         |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*ChannelCreate(), \_msg\_info, MsgInfo(), MsgRead(), MsgReadv(), MsgReceive(), MsgReceivePulse(), MsgReceivePulsev(), MsgReply(), MsgReplyv(), MsgWrite(), MsgWritev(), \_pulse, TimerTimeout()*

## ***MsgReply()*, *MsgReply\_r()***

© 2004, QNX Software Systems Ltd.

*Reply with a message*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgReply(int rvid,
 int status,
 const void* msg,
 int size);

int MsgReply_r(int rvid,
 int status,
 const void* msg,
 int size);
```

### **Arguments:**

|               |                                                                                         |
|---------------|-----------------------------------------------------------------------------------------|
| <i>rvid</i>   | The receive ID that <i>MsgReceive*()</i> returned when you received the message.        |
| <i>status</i> | The status to use when unblocking the <i>MsgSend*()</i> call in the <i>rvid</i> thread. |
| <i>msg</i>    | A pointer to a buffer that contains the message that you want to reply with.            |
| <i>size</i>   | The size of the message, in bytes.                                                      |

### **Library:**

`libc`

### **Description:**

The *MsgReply()* and *MsgReply\_r()* kernel calls reply with a message to the thread identified by *rvid*. The thread being replied to must be in the REPLY-blocked state. Any thread in the receiving process is free to reply to the message, however, it may be replied to only once for each receive.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The *MsgSend\*()* in the *rvid* thread unblocks with a return value of *status*.



---

The status argument for both *MsgReply()*, and *MsgReply\_r()* functions should not be passed a negative value.

---

The number of bytes transferred is the minimum of that specified by both the replier and the sender. The reply data isn't allowed to overflow the reply buffer area provided by the sender.

The data transfer occurs immediately, and the replying task doesn't block. There's no need to reply to received messages in any particular order, but you must eventually reply to each message to allow the sending thread(s) to continue execution.

### Blocking states

None. In the network case, lower priority threads may run.

### Native networking

The *MsgReply()* function has increased latency when it's used to communicate across a network — the server is now writing data to its local **npm-qnet**, which may need to communicate with the client's **npm-qnet** to actually transfer the data.

### Returns:

The only difference between the *MsgReply()* and *MsgReply\_r()* functions is the way they indicate errors:

*MsgReply()*      If an error occurs, -1 is returned and *errno* is set.

*MsgReply\_r()*    This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

|            |                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffers.                |
| ESRCH      | The thread indicated by <i>rcvid</i> doesn't exist, or is no longer REPLY-blocked on the channel, or the connection is detached. |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                                           |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*MsgReceive(), MsgReceivev(), MsgReplyv(), MsgSend(), MsgSendv(), MsgWrite(), MsgWritev()*



**Synopsis:**

```
#include <sys/neutrino.h>

int MsgReplyv(int rvid,
 int status,
 const iov_t* riov,
 int rparts);

int MsgReplyv_r(int rvid,
 int status,
 const iov_t* riov,
 int rparts);
```

**Arguments:**

*rvid*      The receive ID that *MsgReceive\*()* returned when you received the message.

*status*    The status to use when unblocking the *MsgSend\*()* call in the *rvid* thread.

*riov*      An array of buffers that contains the message that you want to reply with.

*size*      The number of elements in the array.

**Library:**

`libc`

**Description:**

The *MsgReplyv()* and *MsgReplyv\_r()* kernel calls reply with a message to the thread identified by *rvid*. The thread being replied to must be in the REPLY-blocked state. Any thread in the receiving process is free to reply to the message, however, it may be replied to only once for each receive.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The *MsgSend\*()* in the *rvid* thread unblocks with a return value of *status*.

The data is taken from the array of message buffers pointed to by *riov*. The number of elements in this array is given by *rparts*. The size of the message is the sum of the sizes of each buffer.

The number of bytes transferred is the minimum of that specified by both the replier and the sender. The reply data isn't allowed to overflow the reply buffer area provided by the sender.

The data transfer occurs immediately, and the replying task doesn't block. There's no need to reply to received messages in any particular order, but you must eventually reply to each message to allow the sending thread(s) to continue execution.

It's quite common to reply with two-part messages consisting of a fixed header and a buffer of data. The *MsgReplyv()* function gathers the data from the buffer list into a logical contiguous message and transfers it to the sender's reply buffer(s). The sender doesn't need to specify the same number or size of buffers. The data is laid down filling each buffer as required. The filesystem, for example, builds a reply list pointing into its cache in order to reply with what appears to be one contiguous piece of data.

### **Blocking states**

None. In the network case, lower priority threads may run.

### **Returns:**

The only difference between the *MsgReplyv()* and *MsgReplyv\_r()* functions is the way they indicate errors:

*MsgReplyv()* If an error occurs, -1 is returned and *errno* is set.

*MsgReplyv\_r()* This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

**Errors:**

|            |                                                                                                                                  |
|------------|----------------------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffers.                |
| ESRCH      | The thread indicated by <i>rcvid</i> doesn't exist, or is no longer REPLY-blocked on the channel, or the connection is detached. |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                                           |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgReceive(), MsgReceivev(), MsgReply(), MsgSend(), MsgSendv(), MsgWrite(), MsgWritev()*

# ***MsgSend(), MsgSend\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSend(int coid,
 const void* msg,
 int sbytes,
 void* rmsg,
 int rbytes);

int MsgSend_r(int coid,
 const void* msg,
 int sbytes,
 void* rmsg,
 int rbytes);
```

## **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*msg*      A pointer to a buffer that contains the message that you want to send.

*sbytes*    The number of bytes to send.

*rmsg*      A pointer to a buffer where the reply can be stored.

*rbytes*    The size of the reply buffer, in bytes.

## **Library:**

libc

## **Description:**

The *MsgSend()* and *MsgSend\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSend()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Native networking

When a client sends a message to a remote server, the client is effectively sending the message via its local microkernel; the network manager does the actual "work." The local network manager negotiates with the remote network manager and causes the message to be delivered there. However, the remote manager is the one that actually delivers the message to the server.

This message transfer from the remote manager to the server is accomplished via a special nonblocking message pass.

The only impact on the client is the latency of the message-passing operations. This is purely a function of the network link speed and the overhead associated with the protocol (i.e. **npm-qnet** for native networking) that **io-net** uses.

The client still remains blocked in its *MsgSend()*, and unblocks only on account of a signal, a kernel timeout, or the completion of its function.

## Returns:

The only difference between the *MsgSend()* and *MsgSend\_r()* functions is the way they indicate errors:

|                    |                |                                                                                                                                                                                                                                                                          |
|--------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSend()</i>   | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                    | -1             | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ).                                                                                                                |
| <i>MsgSend_r()</i> | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                    | negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

## Errors:

|       |                                                                                                                                                                                                                                                              |
|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|            |                                                                                                                            |
|------------|----------------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply. |
| EINTR      | The call was interrupted by a signal.                                                                                      |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                  |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                           |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendsvnc(), MsgSendv(), MsgSendvnc(), MsgSendvs(), MsgSendvsnc(), TimerTimeout()*

## ***MsgSendnc()*, *MsgSendnc\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel (non-cancellation point)*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendnc(int coid,
 const void* msg,
 int sbytes,
 void* rmsg,
 int rbytes);

int MsgSendnc_r(int coid,
 const void* msg,
 int sbytes,
 void* rmsg,
 int rbytes);
```

### **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*msg*      A pointer to a buffer that contains the message that you want to send.

*sbytes*    The number of bytes to send.

*rmsg*      A pointer to a buffer where the reply can be stored.

*rbytes*    The size of the reply buffer, in bytes.

### **Library:**

libc

### **Description:**

The *MsgSendnc()* and *MsgSendnc\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.



The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSend()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendnc()* and *MsgSendnc\_r()* functions is the way they indicate errors:

|                    |         |                                                                                                                                                           |
|--------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendnc()</i> | Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
|                    | -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

|                      |                |                                                                                                                                                                                                                                                                          |
|----------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendnc_r()</i> | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                      | negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

## Errors:

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                             |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendPulse(), MsgSendsv(), MsgSendsvnc(), MsgSendv(),  
MsgSendvnc(), MsgSendvs(), MsgSendvsnc(), TimerTimeout()*

## ***MsgSendPulse()*, *MsgSendPulse\_r()*** © 2004, QNX Software Systems

Ltd.

*Send a pulse to a process*

---

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendPulse (int coid,
 int priority,
 int code,
 int value);

int MsgSendPulse_r (int coid,
 int priority,
 int code,
 int value);
```

### **Arguments:**

- coid*            The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.
- priority*        The priority to use for the pulse. This must be within the range of valid priorities, which you can determine by calling *sched\_get\_priority\_min()* and *sched\_get\_priority\_max()*.
- code*            The 8-bit pulse code.
- Although *code* can be any 8-bit signed value, you should avoid *code* values less than zero, in order to avoid conflict with kernel- or QNX manager-generated pulse codes. These codes all start with `_PULSE_CODE_` and are defined in `<sys/neutrino.h>`; for more information, see the documentation for the `_pulse` structure. A safe range of pulse values is `_PULSE_CODE_MINAVAIL` through `_PULSE_CODE_MAXAVAIL`.
- value*           The 32-bit pulse value.

## Library:

`libc`

## Description:

The *MsgSendPulse()* and *MsgSendPulse\_r()* kernel calls send a short, nonblocking message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

You can send a pulse to a process if the sending process's real or effective user ID either:

- matches the real or effective user ID of the receiving process

Or:

- equals zero.

This permission checking is identical to that used by *kill()*.

You can use *MsgSendPulse()* for many purposes; however, due to the small payload of data, you shouldn't use it for transmitting large amounts of bulk data by sending a great number of pulses.

Pulses are queued for the receiving process in the system, using a dynamic pool of memory objects. If pulses are generated faster than they can be consumed by the receiver, then over a period of time the system queue for the pulses could reach a low memory condition. If there's no memory available for the pulse to be queued in the system, the kernel fails the pulse request with an error of *EAGAIN*. If the priority, code and value don't change, the kernel compresses the pulses by storing an 8-bit count with an already queued pulse.

When you receive a pulse via the *MsgReceive\*()* kernel call, the *rcvid* returned is zero. This indicates to the receiver that it's a pulse and, unlike a message, shouldn't be replied to using *MsgReply\*()*.



In a client/server design, *MsgDeliverEvent()* is typically used in the server, and *MsgSendPulse()* in the client.

---

## Blocking states

None. In the network case, lower priority threads may run.

## Native networking

You can use *MsgSendPulse()* to send pulses across the network.

## Returns:

The only difference between the *MsgSendPulse()* and *MsgSendPulse\_r()* functions is the way they indicate errors:

### *MsgSendPulse()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

### *MsgSendPulse\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

|        |                                                                                                                                                                                                                                                            |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The kernel had insufficient resources to enqueue the pulse.                                                                                                                                                                                                |
| EBADF  | The connection indicated by <i>coid</i> is no longer connected to a channel or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server or the network manager if it failed to respond to multiple polls. |
| EFAULT | A fault occurred when the kernel tried to access the buffers provided.                                                                                                                                                                                     |

|            |                                                                                                                                 |
|------------|---------------------------------------------------------------------------------------------------------------------------------|
| EPERM      | This process doesn't have sufficient permission to send a pulse to the connection, <i>coid</i> .                                |
| ESRVRFAULT | A fault occurred in the server's address space when it tried to write the pulse message to the server's receive message buffer. |

### Classification:

QNX Neutrino

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### Caveats:

If the server faults on delivery, the pulse is either lost or an error is returned.

### See also:

*MsgDeliverEvent(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendsv(), MsgSendsvnc(), MsgSendv(),  
MsgSendvnc(), MsgSendvs(), MsgSendvsnc(), **\_pulse**,  
sched\_get\_priority\_min(), sched\_get\_priority\_max()*

# ***MsgSendsv()*, *MsgSendsv\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendsv(int coid,
 const void* smsg,
 int sbytes,
 const iov_t* riov,
 int rparts);

int MsgSendsv_r(int coid,
 const void* smsg,
 int sbytes,
 const iov_t* riov,
 int rparts);
```

## **Arguments:**

|               |                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------|
| <i>coid</i>   | The ID of the channel to send the message on, which you've established by calling <i>ConnectAttach()</i> . |
| <i>smsg</i>   | A pointer to a buffer that contains the message that you want to send.                                     |
| <i>sbytes</i> | The number of bytes to send.                                                                               |
| <i>riov</i>   | An array of buffers where the reply can be stored.                                                         |
| <i>rparts</i> | The number of elements in the <i>riov</i> array.                                                           |

## **Library:**

libc

## **Description:**

The *MsgSendsv()* and *MsgSendsv\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.



The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSendsv()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendsvnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendsv()* and *MsgSendsv\_r()* functions is the way they indicate errors:

|                    |         |                                                                                                                                                           |
|--------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendsv()</i> | Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
|                    | -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

|                      |                |                                                                                                                                                                                                                                                                          |
|----------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendsv_r()</i> | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                      | negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

## Errors:

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                             |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendPulse(), MsgSendsvnc(), MsgSendv(),  
MsgSendvnc(), MsgSendvs(), MsgSendvsnc(), TimerTimeout()*

## ***MsgSendsvnc()*, *MsgSendsvnc\_r()*** © 2004, QNX Software Systems Ltd.

*Send a message to a channel (non-cancellation point)*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendsvnc(int coid,
 const void* msg,
 int sbytes,
 const iov_t* riov,
 int rparts);

int MsgSendsvnc_r(int coid,
 const void* msg,
 int sbytes,
 const iov_t* riov,
 int rparts);
```

### **Arguments:**

|               |                                                                                                            |
|---------------|------------------------------------------------------------------------------------------------------------|
| <i>coid</i>   | The ID of the channel to send the message on, which you've established by calling <i>ConnectAttach()</i> . |
| <i>msg</i>    | A pointer to a buffer that contains the message that you want to send.                                     |
| <i>sbytes</i> | The number of bytes to send.                                                                               |
| <i>riov</i>   | An array of buffers where the reply can be stored.                                                         |
| <i>rparts</i> | The number of elements in the <i>riov</i> array.                                                           |

### **Library:**

libc

### **Description:**

The *MsgSendsvnc()* and *MsgSendsvnc\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSendsv()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendsvnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendsvnc()* and *MsgSendsvnc\_r()* functions is the way they indicate errors:

#### *MsgSendsvnc()*

|         |                                                                                                                                                           |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
| -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

*MsgSendsvnc\_r()*

|                |                                                                                                                                                                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success        | The value of <i>status</i> from <i>MsgReply*</i> ().                                                                                                                                                                                                                   |
| negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError</i> ()). |

**Errors:**

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout</i> ().                                                                                                                                                                                             |

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendv(),  
MsgSendsvnc(), MsgSendsvs(), MsgSendsvsnc(), TimerTimeout()*

# ***MsgSendv()*, *MsgSendv\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendv(int coid,
 const iov_t* siov,
 int sparts,
 const iov_t* riov,
 int rparts);

int MsgSendv_r(int coid,
 const iov_t* siov,
 int sparts,
 const iov_t* riov,
 int rparts);
```

## **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*siov*      An array of buffers that contains the message that you want to send.

*sparts*    The number of elements in the *siov* array.

*riov*      An array of buffers where the reply can be stored.

*rparts*    The number of elements in the *riov* array.

## **Library:**

`libc`

## **Description:**

The *MsgSendv()* and *MsgSendv\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.



The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

It's quite common to send two-part messages consisting of a fixed header and a buffer of data. The *MsgSendv()* function gathers the data from the send list into a logically contiguous message and transfers it to the receiver. The receiver doesn't need to specify the same number or size of buffers. The data is laid down filling each entry as required. The same applies to the replied data.

*MsgSendv()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendvnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

## Returns:

The only difference between the *MsgSendv()* and *MsgSendv\_r()* functions is the way they indicate errors:

|                     |                |                                                                                                                                                                                                                                                                          |
|---------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendv()</i>   | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                     | -1             | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ).                                                                                                                |
| <i>MsgSendv_r()</i> | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                     | negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

## Errors:

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |

ETIMEDOUT      A kernel timeout unblocked the call. See *TimerTimeout()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(), MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendsvnc(), MsgSendvnc(), MsgSendvs(), MsgSendvsnc(), TimerTimeout()*

## ***MsgSendvnc()*, *MsgSendvnc\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel (non-cancellation point)*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendvnc(int coid,
 const iov_t* siov,
 int sparts,
 const iov_t* riov,
 int rparts);

int MsgSendvnc_r(int coid,
 const iov_t* siov,
 int sparts,
 const iov_t* riov,
 int rparts);
```

### **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*siov*      An array of buffers that contains the message that you want to send.

*sparts*    The number of elements in the *siov* array.

*riov*      An array of buffers where the reply can be stored.

*rparts*    The number of elements in the *riov* array.

### **Library:**

libc

### **Description:**

The *MsgSendvnc()* and *MsgSendvnc\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSendv()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendvnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendvnc()* and *MsgSendvnc\_r()* functions is the way they indicate errors:

#### *MsgSendvnc()*

|         |                                                                                                                                                           |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
| -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

*MsgSendvnc\_r()*

|                |                                                                                                                                                                                                                                                                        |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success        | The value of <i>status</i> from <i>MsgReply*</i> ().                                                                                                                                                                                                                   |
| negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError</i> ()). |

**Errors:**

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout</i> ().                                                                                                                                                                                             |

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendsvnc(),  
MsgSendv(), MsgSendvs(), MsgSendvsnc(), TimerTimeout()*

# ***MsgSendvs()*, *MsgSendvs\_r()***

© 2004, QNX Software Systems Ltd.

*Send a message to a channel*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendvs(int coid,
 const iov_t* siov,
 int sparts,
 void* rmsg,
 int rbytes);

int MsgSendvs_r(int coid,
 const iov_t* siov,
 int sparts,
 void* rmsg,
 int rbytes);
```

## **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*siov*      An array of buffers that contains the message that you want to send.

*sparts*    The number of elements in the *siov* array.

*rmsg*      A pointer to a buffer where the reply can be stored.

*rbytes*    The size of the reply buffer.

## **Library:**

libc

## **Description:**

The *MsgSendvs()* and *MsgSendvs\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.



The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSendvs()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendvsnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendvs()* and *MsgSendvs\_r()* functions is the way they indicate errors:

|                    |         |                                                                                                                                                           |
|--------------------|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendvs()</i> | Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
|                    | -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

|                      |                |                                                                                                                                                                                                                                                                          |
|----------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgSendvs_r()</i> | Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
|                      | negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

## Errors:

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                             |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendsvnc(),  
MsgSendv(), MsgSendvnc(), MsgSendvsnc(), TimerTimeout()*

## ***MsgSendvsnc()*, *MsgSendvsnc\_r()*** © 2004, QNX Software Systems Ltd.

*Send a message to a channel (non-cancellation point)*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgSendvsnc(int coid,
 const iov_t* siov,
 int sparts,
 void* rmsg,
 int rbytes);

int MsgSendvsnc_r(int coid,
 const iov_t* siov,
 int sparts,
 void* rmsg,
 int rbytes);
```

### **Arguments:**

*coid*      The ID of the channel to send the message on, which you've established by calling *ConnectAttach()*.

*siov*      An array of buffers that contains the message that you want to send.

*sparts*    The number of elements in the *siov* array.

*rmsg*      A pointer to a buffer where the reply can be stored.

*rbytes*    The size of the reply buffer.

### **Library:**

`libc`

### **Description:**

The *MsgSendvsnc()* and *MsgSendvsnc\_r()* kernel calls send a message to a process's channel identified by *coid*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The number of bytes transferred is the minimum of that specified by both the sender and the receiver. The send data isn't allowed to overflow the receive buffer area provided by the receiver. The reply data isn't allowed to overflow the reply buffer area provided.

The sending thread becomes blocked waiting for a reply. If the receiving process has a thread that's RECEIVE-blocked on the channel, the transfer of data into its address space occurs immediately, and the receiving thread is unblocked and made ready to run. The sending thread becomes REPLY-blocked. If there are no waiting threads on the channel, the sending thread becomes SEND-blocked and is placed in a queue (perhaps with other threads). In this case, the actual transfer of data doesn't occur until a receiving thread receives on the channel. At this point, the sending thread becomes REPLY-blocked.

*MsgSendvs()* is a cancellation point for the *ThreadCancel()* kernel call; *MsgSendvsnc()* isn't.

### Blocking states

|             |                                                                                                                                                                           |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| STATE_SEND  | The message has been sent but not yet received. If a thread is waiting to receive the message, this state is skipped and the calling thread goes directly to STATE_REPLY. |
| STATE_REPLY | The message has been received but not yet replied to. This state may be entered directly, or from STATE_SEND.                                                             |

### Returns:

The only difference between the *MsgSendvsnc()* and *MsgSendvsnc\_r()* functions is the way they indicate errors:

#### *MsgSendvsnc()*

|         |                                                                                                                                                           |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                      |
| -1      | An error occurred ( <i>errno</i> is set), or the server called <i>MsgError*()</i> ( <i>errno</i> is set to the error value passed to <i>MsgError()</i> ). |

*MsgSendvsnc\_r()*

|                |                                                                                                                                                                                                                                                                          |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Success        | The value of <i>status</i> from <i>MsgReply*()</i> .                                                                                                                                                                                                                     |
| negative value | An error occurred ( <i>errno</i> is <i>NOT</i> set, the value is the negative of a value from the Errors section), or the server called <i>MsgError*()</i> ( <i>errno</i> is <i>NOT</i> set, the value is the negative of the error value passed to <i>MsgError()</i> ). |

**Errors:**

|            |                                                                                                                                                                                                                                                              |
|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBADF      | The connection indicated by <i>coid</i> is no longer connected to a channel, or the connection indicated by <i>coid</i> doesn't exist. The channel may have been terminated by the server, or the network manager if it failed to respond to multiple polls. |
| EFAULT     | A fault occurred when the kernel tried to access the buffers provided. This may have occurred on the receive or the reply.                                                                                                                                   |
| EINTR      | The call was interrupted by a signal.                                                                                                                                                                                                                        |
| ESRCH      | The server died while the calling thread was SEND-blocked or REPLY-blocked.                                                                                                                                                                                  |
| ESRVRFAULT | A fault occurred in a server's address space when accessing the server's message buffers.                                                                                                                                                                    |
| ETIMEDOUT  | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                                                                                                                                                                             |

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*ConnectAttach(), MsgReceive(), MsgReceivev(), MsgSend(),  
MsgSendnc(), MsgSendPulse(), MsgSendsv(), MsgSendsvnc(),  
MsgSendv(), MsgSendvnc(), MsgSendvs(), TimerTimeout()*

# ***MsgVerifyEvent()*, *MsgVerifyEvent\_r()***

© 2004, QNX Software

Systems Ltd.

*Check the validity of a receive ID and an event configuration*

---

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgVerifyEvent(int rvid,
 const struct sigevent event);

int MsgVerifyEvent_r(int rvid,
 const struct sigevent event);
```

## **Arguments:**

*rvid*     The receive ID that you want to check.

*event*    A pointer to a **sigevent** structure that contains the event you want to check.

## **Library:**

**libc**

## **Description:**

The *MsgVerifyEvent()* and *MsgVerifyEvent\_r()* kernel calls check the validity of the receive ID *rvid*, and the *event* configuration. You can use these functions to verify that an event is well-formed by a client (pass a *rvid* of 0), and by a server (pass a *rvid* of the target thread).

These functions are identical except in the way they indicate errors. See the Returns section for details.

## **Blocking states**

These calls don't block.

## **Returns:**

The only difference between the *MsgVerifyEvent()* and *MsgVerifyEvent\_r()* functions is the way they indicate errors:



*MsgVerifyEvent()*

If an error occurs, -1 is returned and *errno* is set.

*MsgVerifyEvent\_r()*

This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

**Errors:**

- EBADF      The channel for the pulse delivery doesn't exist.
- EINVAL     Invalid *event* structure.
- ESRCH      The connection for the pulse doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgReceive()*, *MsgReceivev()*, *MsgReply()*, *MsgSend()*, *MsgSendv()*,  
*MsgWrite()*, *MsgWritev()*, **sigevent**

## ***MsgWrite()*, *MsgWrite\_r()***

© 2004, QNX Software Systems Ltd.

*Write a reply*

### **Synopsis:**

```
#include <sys/neutrino.h>

int MsgWrite(int rvid,
 const void* msg,
 int size,
 int offset);

int MsgWrite_r(int rvid,
 const void* msg,
 int size,
 int offset);
```

### **Arguments:**

*rvid*     The value returned by *MsgReceive\*()* when you received the message.

*msg*     A pointer to a buffer that contains the data you want to write.

*size*     The number of bytes that you want to write. These functions don't let you write past the end of the sender's buffer; they return the number of bytes actually written.

*offset*   An offset into the sender's buffer that indicates where you want to start writing the data.

### **Library:**

`libc`

### **Description:**

The *MsgWrite()* and *MsgWrite\_r()* kernel calls write data to the reply buffer of a thread identified by *rvid*. The thread being written to must be in the REPLY-blocked state. Any thread in the receiving process is free to write to the reply message.

These functions are identical except in the way they indicate errors. See the Returns section for details.

You use this function in one of these situations:

- The data arrives over time and is quite large. Rather than buffer all the data, you can use *MsgWrite()* to write it into the destination thread's reply message buffer, as it arrives.
- Messages are received that are larger than available buffer space. Perhaps the process is an agent between two processes and simply filters the data and passes it on. You can use *MsgRead\*()* to read messages in small pieces, and use *MsgWrite()* to write messages in small pieces.

To complete a message exchange, you must call *MsgReply\*()*. The reply doesn't need to contain any data. If it does contain data, then the data is always written at offset zero in the destination thread's reply message buffer. This is a convenient way of writing the header once all of the information has been gathered.

A single call to *MsgReply\*()* is always more efficient than calls to *MsgWrite()* followed by a call to *MsgReply\*()*.

### Blocking states

None. In the network case, lower priority threads may run.

### Native networking

The *MsgWrite()* function has increased latency when you use it to communicate across a network — the server is now writing data to its local **npm-qnet**, which may need to communicate with the client's **npm-qnet** to actually transfer the data. The server's *MsgWrite()* call effectively sends a message to the server's **npm-qnet** to initiate this data transfer.

But since the server's **npm-qnet** has no way to determine the size of the client's receive data area, the number of bytes reported as having been transferred by the server during its *MsgWrite()* call *might not be accurate* — the reported number will instead reflect the number of bytes transferred by the server to its **npm-qnet**.

The message is buffered in the server side's `npm-qnet` until the client replies, in order to reduce the number of network transactions.

If you want to determine the size of the sender's reply buffer, set the `_NTO_CHF_REPLY_LEN` when you call *ChannelCreate()*.

## Returns:

The only difference between the *MsgWrite()* and *MsgWrite\_r()* functions is the way they indicate errors:

|                     |                                                                                                                                                                |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>MsgWrite()</i>   | The number of bytes written. If an error occurs, -1 is returned and <i>errno</i> is set.                                                                       |
| <i>MsgWrite_r()</i> | The number of bytes written. This function does <b>NOT</b> set <i>errno</i> . If an error occurs, the negative of a value from the Errors section is returned. |

## Errors:

|            |                                                                                                                  |
|------------|------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffer. |
| ESRCH      | The thread indicated by <i>rvid</i> doesn't exist or its connection was detached.                                |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                           |

## Classification:

QNX Neutrino

### Safety

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*ChannelCreate(), MsgRead(), MsgReadv(), MsgReceive(),  
MsgReceivev(), MsgReply(), MsgReplyv(), MsgWritev()*

# ***MsgWritev()*, *MsgWritev\_r()***

© 2004, QNX Software Systems Ltd.

*Write a reply*

## **Synopsis:**

```
#include <sys/neutrino.h>

int MsgWritev(int rvid,
 const iov_t* iov,
 int parts,
 int offset);

int MsgWritev_r(int rvid,
 const iov_t* iov,
 int parts,
 int offset);
```

## **Arguments:**

*rvid*      The value returned by *MsgReceive\*()* when you received the message.

*iov*        An array of buffers that contains the data you want to write.

*parts*     The number of elements in the array. These functions don't let you write past the end of the sender's buffer; they return the number of bytes actually written.

*offset*    An offset into the sender's buffer that indicates where you want to start writing the data.

## **Library:**

libc

## **Description:**

The *MsgWritev()* and *MsgWritev\_r()* kernel calls write data to the reply buffer of a thread identified by *rvid*. The thread being written to must be in the REPLY-blocked state. Any thread in the receiving process is free to write to the reply message.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The data transfer occurs immediately and your thread doesn't block. The state of the sending thread doesn't change.

You'll use this function in one of these situations:

- The data arrives over time and is quite large. Rather than buffer all the data, you can use *MsgWritev()* to write it into the destination thread's reply message buffer, as it arrives.
- Messages are received that are larger than available buffer space. Perhaps the process is an agent between two processes and simply filters the data and passes it on. You can use *MsgRead\*()* to read messages in small pieces, and use *MsgWritev()* to write messages in small pieces.

To complete a message exchange, you must call *MsgReply\*()*. The reply doesn't need to contain any data. If it does contain data, then the data is always written at offset zero in the destination thread's reply message buffer. This is a convenient way of writing the header once all of the information has been gathered.

A single call to *MsgReply\*()* is always more efficient than calls to *MsgWritev()* followed by a call to *MsgReply\*()*.

### Blocking states

None. In the network case, lower priority threads may run.

### Returns:

The only difference between the *MsgWritev()* and *MsgWritev\_r()* functions is the way they indicate errors:

*MsgWritev()*      The number of bytes written. If an error occurs, -1 is returned and *errno* is set.

*MsgWritev\_r()*    The number of bytes written. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

|            |                                                                                                                  |
|------------|------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffer. |
| ESRCH      | The thread indicated by <i>rcvid</i> doesn't exist or its connection was detached.                               |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                           |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*MsgRead(), MsgReady(), MsgReceive(), MsgReceivev(), MsgReply(), MsgReplyv(), MsgWrite()*



**Synopsis:**

```
#include <sys/mman.h>

int msync(void * addr,
 size_t len,
 int flags);
```

**Arguments:**

- addr*     The beginning of the range of addresses that you want to synchronize.
- len*     The length of the range of addresses, in bytes.
- flags*    A bitwise inclusive OR of one or more of the following flags:
- **MS\_ASYNC** — perform asynchronous writes. The function returns immediately once all the write operations are initiated or queued for servicing.
  - **MS\_INVALIDATE** — invalidate cached data. Invalidates all cached copies of mapped data that are inconsistent with the permanent storage locations such that subsequent references obtain data that was consistent with the permanent storage locations sometime between the call to *msync()* and the first subsequent memory reference to the data.
  - **MS\_INVALIDATE\_ICACHE** — (QNX Neutrino extension) if you're dynamically modifying code, use this flag to make sure that the new code is what will be executed.
  - **MS\_SYNC** — perform synchronous writes. The function doesn't return until all write operations are completed as defined for synchronized I/O data integrity completion.



---

You can specify at most one of MS\_ASYNC and MS\_SYNC, not both.

---

**Library:**

`libc`

**Description:**

The *msync()* function writes all modified data to permanent storage locations, if any, in those whole pages containing any part of the address space of the process starting at address *addr* and continuing for *len* bytes. The *msync()* function is used with memory mapped files. If no such storage exists, *msync()* need not have any effect. If requested, the *msync()* function then invalidates cached copies of data.

For mappings to files, this function ensures that all write operations are completed as defined for synchronized I/O data integrity completion.



---

Mappings to files aren't implemented on all filesystems.

---

When the *msync()* function is called on MAP\_PRIVATE mappings, any modified data won't be written to the underlying object and won't cause such data to be made visible to other processes.

The behavior of *msync()* is unspecified if the mapping wasn't established by a call to *mmap()*.

If *msync()* causes any write to a file, the file's *st\_ctime* and *st\_mtime* fields are marked for update.

**Returns:**

- 0      Success
- 1     An error occurred (*errno* is set).

## Errors:

|        |                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBUSY  | Some or all of the addresses in the range starting at <i>addr</i> and continuing for <i>len</i> bytes are locked, and MS_INVALIDATE is specified.                                                        |
| EINVAL | Invalid <i>flags</i> value.                                                                                                                                                                              |
| ENOMEM | The addresses in the range starting at <i>addr</i> and continuing for <i>len</i> bytes are outside the range allowed for the address space of a process or specify one or more pages that aren't mapped. |

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

MS\_INVALIDATE\_ICACHE is a QNX Neutrino extension.

## See also:

*mmap()*, *sysconf()*

## ***munlock()***

*Unlock a buffer*

© 2004, QNX Software Systems Ltd.

### **Synopsis:**

```
#include <sys/mman.h>

int munlock(const void * addr,
 size_t len);
```

### **Library:**

libc

### **Description:**

The *munlock()* function isn't currently supported.

### **Returns:**

-1 to indicate an error (*errno* is set).

### **Errors:**

ENOSYS      The *munlock()* function isn't currently supported.

### **Classification:**

POSIX 1003.1 (Realtime Extensions)

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mlock(), mlockall(), munlockall()*

# ***munlockall()***

© 2004, QNX Software Systems Ltd.

*Unlock a process's address space*

---

## **Synopsis:**

```
#include <sys/mman.h>

int munlockall(void);
```

## **Library:**

libc

## **Description:**

The current implementation of the *munlockall()* function doesn't do anything.

## **Returns:**

-1 to indicate an error (*errno* is set).

## **Errors:**

ENOSYS     The *munlockall()* function isn't currently supported.

## **Classification:**

POSIX 1003.1 (Realtime Extensions)

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mlock(), munlock(), mlockall()*

# ***munmap()***

© 2004, QNX Software Systems Ltd.

*Unmap previously mapped addresses*

## **Synopsis:**

```
#include <sys/mman.h>

int munmap(void * addr,
 size_t len);
```

## **Arguments:**

*addr*     The beginning of the range of addresses that you want to unmap.

*len*      The length of the range of addresses, in bytes.

## **Library:**

`libc`

## **Description:**

The *munmap()* function removes any mappings for pages in the address range starting at *addr* and continuing for *len* bytes, rounded up to the next multiple of the page size. Subsequent references to these pages cause a SIGSEGV signal to be set on the process.

If there are no mappings in the specified address range, then *munmap()* has no effect.

## **Returns:**

0        Success.

-1       Failure; *errno* is set.

## **Errors:**

EINVAL    The addresses in the specified range are outside the range allowed for the address space of a process.

ENOSYS    The function *munmap()* isn't supported by this implementation.



## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Currently, you can't *munmap()* just a part of an area mapped with *mmap()*.

## See also:

*mmap()*, *mprotect()*, *shm\_open()*, *shm\_unlink()*

## ***munmap\_device\_io()***

© 2004, QNX Software Systems Ltd.

*Free access to a device's registers*

### **Synopsis:**

```
#include <sys/mman.h>

int munmap_device_io(uintptr_t io,
 size_t len);
```

### **Arguments:**

*io*      The address of the area that you want to unmap.

*len*     The number of bytes of device I/O memory that you want to unmap.

### **Library:**

libc

### **Description:**

The function *munmap\_device\_io()* unmaps *len* bytes of device I/O memory at *io* (that was previously mapped with *mmap\_device\_io()*).

### **Returns:**

-1      An error occurred (*errno* is set).

Any other value

Successful unmapping.

### **Errors:**

EINVAL      The addresses in the specified range are outside the range allowed for the address space of a process.

ENOSYS     The function *munmap()* isn't supported by this implementation.

ENXIO      The address from *io* for *len* bytes is invalid.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*mmap\_device\_io()*, *munmap()*

## ***munmap\_device\_memory()***

© 2004, QNX Software Systems Ltd.

*Unmap previously mapped addresses*

### **Synopsis:**

```
#include <sys/mman.h>

int munmap_device_memory(void * addr,
 size_t len);
```

### **Arguments:**

*addr*     The beginning of the range of addresses that you want to unmap.

*len*      The length of the range of addresses, in bytes.

### **Library:**

`libc`

### **Description:**

The *munmap\_device\_memory()* function is essentially the same as *munmap()*. It removes any mappings for pages in the address range starting at *addr* and continuing for *len* bytes, rounded up to the next multiple of the page size. Subsequent references to these pages cause a SIGSEGV signal to be set on the process.

If there are no mappings in the specified address range, then *munmap()* has no effect.

This function is the complement of *mmap\_device\_memory()*.

### **Returns:**

-1     An error occurred (*errno* is set).

Any other value

Success.

## Errors:

- |        |                                                                                                        |
|--------|--------------------------------------------------------------------------------------------------------|
| EINVAL | The addresses in the specified range are outside the range allowed for the address space of a process. |
| ENOSYS | The <i>munmap()</i> function isn't supported by this implementation.                                   |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*mmap\_device\_memory()*, *munmap()*, *munmap\_device\_io()*

## ***name\_attach()***

© 2004, QNX Software Systems Ltd.

*Register a name in the namespace and create a channel*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

name_attach_t * name_attach(dispatch_t * dpp,
 const char * path,
 unsigned flags);
```

### **Arguments:**

*dpp* NULL, or a dispatch handle returned by *dispatch\_create()*.

*path* The path that you want to register under */dev/name/[local|global]/*. This name shouldn't contain any *..* characters or start with a leading slash */*.

*flags* Flags that affect the function's behavior:

- NAME\_FLAG\_ATTACH\_GLOBAL — attach the name globally instead of locally.

### **Library:**

libc

### **Description:**

The *name\_attach()*, *name\_close()*, *name\_detach()*, and *name\_open()* functions provide the basic pathname-to-server-connection mapping, without having to become a *full* resource manager.

A dispatch structure is created for you automatically if you pass NULL as the *dpp*. If you've already created a dispatch structure, pass it in as the *dpp*. If you provide your own *dpp*, set *flags* to NAME\_FLAG\_DETACH\_SAVEDPP when calling *name\_detach()*; otherwise, your *dpp* is detached and destroyed automatically.

The *name\_attach()* function puts the name *path* into the path namespace under */dev/name/[local|global]/path*. The name is attached locally by default, or globally when you set

NAME\_FLAG\_ATTACH\_GLOBAL in the *flags*. You can see attached names in `/dev/name/local` and `/dev/name/global` directories.

*ChannelCreate()* is called with the `_NTO_CHF_UNBLOCK`, `_NTO_CHF_DISCONNECT`, and `_NTO_CHF_COID_DISCONNECT` flags set. The `_NTO_CHF_THREAD_DEATH` flag isn't passed to *ChannelCreate()*, but is implied by the setting of `_NTO_CHF_COID_DISCONNECT`. Therefore, your server that's using *name\_attach()* may receive pulses as described in *ChannelCreate*. For instance, since `_NTO_CHF_DISCONNECT` is set, when a client calls *name\_close()*, you'll receive the `_PULSE_CODE_DISCONNECT` pulse message.

The *name\_attach()* also receives `_JO_CONNECT` message when *name\_open()* is called.

For more information on the pulses (a *rvid* of 0) related to the `_NTO_CHF_COID_DISCONNECT`, `_NTO_CHF_DISCONNECT`, `_NTO_CHF_THREAD_DEATH` and the `_NTO_CHF_UNBLOCK` flags, see *ChannelCreate()*.

If the receive buffer that the server provides isn't large enough to hold a pulse, then *MsgReceive()* returns -1 with *errno* set to `EFAULT`.

## **name\_attach\_t**

The *name\_attach()* function returns a pointer to a **name\_attach\_t** structure that looks like this:

```
typedef struct _name_attach {
 dispatch_t* dpp;
 int chid;
 int mntid;
 int zero[2];
} name_attach_t;
```

The members include:

- dpp*        The dispatch handle used in the creation of this connection.
- chid*      The channel ID used for *MsgReceive()* directly.
- mntid*     the mount ID for this name.

The information that's generally required by a server using these services is the *chid*.

## Returns:

A pointer to a filled-in `name_attach_t` structure, or NULL if the call fails (*errno* is set).

## Errors:

- EINVAL      Invalid arguments (i.e. a NULL or empty path, a path starts with a leading slash / or contains .. characters).
- ENOMEM     Not enough free memory to complete the operation.
- ENOTDIR    A component of the pathname wasn't a directory entry.

## Examples:

```
#include <stdio.h>
#include <errno.h>
#include <stdlib.h>
#include <sys/dispatch.h>

#define ATTACH_POINT "myname"

/* We specify the header as being at least a pulse */
typedef struct _pulse msg_header_t;

/* Our real data comes after the header */
typedef struct _my_data {
 msg_header_t hdr;
 int data;
} my_data_t;

/** Server Side of the code */
int server() {
 name_attach_t *attach;
 my_data_t msg;
 int rcvid;

 /* Create a local name (/dev/name/local/...) */
 if ((attach = name_attach(NULL, ATTACH_POINT, 0)) == NULL) {
 return EXIT_FAILURE;
 }

 /* Do your MsgReceive's here now with the chid */
```



```

while (1) {
 rcvid = MsgReceive(attach->chid, &msg, sizeof(msg), NULL);

 if (rcvid == -1) { /* Error condition, exit */
 break;
 }

 if (rcvid == 0) { /* Pulse received */
 switch (msg.hdr.code) {
 case _PULSE_CODE_DISCONNECT:
 /*
 * A client disconnected all its connections (called
 * name_close() for each name_open() of our name) or
 * terminated
 */
 ConnectDetach(msg.hdr.scoId);
 break;
 case _PULSE_CODE_UNBLOCK:
 /*
 * REPLY blocked client wants to unblock (was hit by
 * a signal or timed out). It's up to you if you
 * reply now or later.
 */
 break;
 default:
 /*
 * A pulse sent by one of your processes or a
 * _PULSE_CODE_COIDDEATH or _PULSE_CODE_THREADDEATH
 * from the kernel?
 */
 }
 continue;
 }

 /* A QNX IO message received, reject */
 if (msg.hdr.type >= _IO_BASE && msg.hdr.type <= _IO_MAX) {
 MsgError(rcvid, ENOSYS);
 continue;
 }

 /* A message (presumably ours) received, handle */
 printf("Server receive %d \n", msg.data);
 MsgReply(rcvid, EOK, 0, 0);
 }

 /* Remove the name from the space */
 name_detach(attach, 0);

 return EXIT_SUCCESS;
}

```

```
 }

 /** Client Side of the code **/
 int client() {
 my_data_t msg;
 int fd;

 if ((fd = name_open(ATTACH_POINT, 0)) == -1) {
 return EXIT_FAILURE;
 }

 /* We would have pre-defined data to stuff here */
 msg.hdr.type = 0x00;
 msg.hdr.subtype = 0x00;

 /* Do whatever work you wanted with server connection */
 for (msg.data=0; msg.data < 5; msg.data++) {
 printf("Client sending %d \n", msg.data);
 if (MsgSend(fd, &msg, sizeof(msg), NULL, 0) == -1) {
 break;
 }
 }

 /* Close the connection */
 name_close(fd);
 return EXIT_SUCCESS;
 }

 int main(int argc, char **argv) {
 int ret;

 if (argc < 2) {
 printf("Usage %s -s | -c \n", argv[0]);
 ret = EXIT_FAILURE;
 }
 else if (strcmp(argv[1], "-c") == 0) {
 printf("Running Client ... \n");
 ret = client(); /* see name_open() for this code */
 }
 else if (strcmp(argv[1], "-s") == 0) {
 printf("Running Server ... \n");
 ret = server(); /* see name_attach() for this code */
 }
 else {
 printf("Usage %s -s | -c \n", argv[0]);
 ret = EXIT_FAILURE;
 }
 return ret;
 }
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

As a server, you shouldn't assume that you're doing a *MsgReceive()* on a clean channel. In QNX Neutrino (and QNX 4), anyone can create a random message and send it to a process or a channel.

We recommend that you do the following to assure that you're playing safely with others in the system:

```
#include <sys/neutrino.h>

/* All of your messages should start with this header */
typedef struct _pulse msg_header_t;

/* Now your real data comes after this */
typedef struct _my_data {
 msg_header_t hdr;
 int data;
} my_data_t;
```

where:

*hdr*      Contains a *type/subtype* field as the first 4 bytes. This allows you to identify data which isn't destined for your server.

*data* Specifies the receive data structure. The structure must be large enough to contain at least a pulse (which conveniently starts with the *type/subtype* field of most normal messages), because you'll receive a disconnect pulse when clients are detached.

### **See also:**

*ChannelCreate()*, *dispatch\*()* functions, *MsgReceive()*,  
*name\_detach()*, *name\_open()*, *name\_close()*, **\_pulse**, *resmgr\_attach()*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int name_close(int filedes);
```

### **Arguments:**

*filedes*     The file descriptor returned by *name\_open()*.

### **Library:**

`libc`

### **Description:**

The *name\_close()* function closes the *filedes* obtained with the *name\_open()* call.

### **Returns:**

Zero for success, or -1 if an error occurs (*errno* is set).

### **Errors:**

EBADF     Invalid file descriptor *filedes*.

EINTR     The *name\_close()* call was interrupted by a signal.

ENOSYS    The *name\_close()* function isn't implemented for the filesystem specified by *filedes*.

### **Examples:**

See the "Client side of the code" section in *name\_attach()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close()*, *ConnectDetach()*, *name\_attach()*, *name\_detach()*,  
*name\_open()*

*Remove a name from the namespace and destroy the channel*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int name_detach(name_attach_t * attach,
 unsigned flags);
```

### **Arguments:**

- attach*     A pointer to the **name\_attach\_t** structure returned by *name\_attach()*.
- flags*     Flags that affect the function's behavior:
- NAME\_FLAG\_DETACH\_SAVEDPP — don't destroy the dispatch handle.

### **Library:**

libc

### **Description:**

The *name\_detach()* function removes the name from the namespace and destroys the channel created by *name\_attach()*. If you set NAME\_FLAG\_DETACH\_SAVEDPP in *flags*, the dispatch pointer contained in the **name\_attach\_t** structure isn't destroyed; it's up to you to destroy it by calling *dispatch\_destroy()*. The default is to destroy the dispatch pointer.

### **Returns:**

Zero on success, or -1 if an error occurs (*errno* is set).

### **Errors:**

- EINVAL     The mount ID (*mntid*) was never attached with *name\_attach()*.

**Examples:**

See *name\_attach()* and *resmgr\_detach()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ChannelDestroy()*, *dispatch\*()* functions, *name\_attach()*,  
*name\_close()*, *name\_open()*, *resmgr\_detach()*



## Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int name_open(const char * name,
 int flags);
```

## Arguments:

- name*     The name that you want to open for a server connection.
- flags*     Flags that affect the function's behavior:
- NAME\_FLAG\_ATTACH\_GLOBAL — attach the name globally instead of locally.

## Library:

libc

## Description:

The *name\_open()* function opens *name* for a server connection. No ordering is guaranteed when accessing resources on other nodes.



---

Before, when an application used to call *name\_open()* to connect to a service, the server was not aware of that. This has been changed now — a `_IO_CONNECT/_IO_OPEN` message is actually sent to the server.

The server application has to be modified to handle a possible `_IO_CONNECT` message coming in.

For more information, see the technote *Configuring the Global Name Service Manager*.

---

**Returns:**

A nonnegative integer representing a side-channel connection ID (see *ConnectAttach()*) or -1 if an error occurred (*errno* is set).

**Errors:**

|              |                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied on a component of the <i>name</i> .                                                                                                                                                                                                                                                                                                                                         |
| EBADFSYS     | While attempting to open the named file, either the file itself or a component of the path prefix was found to be corrupted. A system failure — from which no automatic recovery is possible — occurred while the file was being written to, or while the directory was being updated. You'll need to invoke appropriate systems-administration procedures to correct this situation before proceeding. |
| EBUSY        | The connection specified by <i>name</i> has already been opened and additional connections aren't permitted.                                                                                                                                                                                                                                                                                            |
| EINTR        | The <i>name_open()</i> operation was interrupted by a signal.                                                                                                                                                                                                                                                                                                                                           |
| EISDIR       | The named path is a directory.                                                                                                                                                                                                                                                                                                                                                                          |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                                                                                                                                                                                                                                          |
| EMFILE       | Too many file descriptors are currently in use by this process.                                                                                                                                                                                                                                                                                                                                         |
| ENAMETOOLONG | The length of the <i>name</i> string exceeds <code>PATH_MAX</code> , or a pathname component is longer than <code>NAME_MAX</code> .                                                                                                                                                                                                                                                                     |
| ENFILE       | Too many files are currently open in the system.                                                                                                                                                                                                                                                                                                                                                        |
| ENOENT       | The connection specified by <i>name</i> doesn't exist.                                                                                                                                                                                                                                                                                                                                                  |
| ENOTDIR      | A component of the <i>name</i> prefix isn't a directory.                                                                                                                                                                                                                                                                                                                                                |

**Examples:**

See *name\_attach()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ConnectAttach()*, *name\_attach()*, *name\_detach()*, *name\_close()*, *open()*

## ***nanosleep()***

© 2004, QNX Software Systems Ltd.

*Suspend a thread until a timeout or signal occurs*

### **Synopsis:**

```
#include <time.h>

int nanosleep(const struct timespec* rntp,
 struct timespec* rmtp);
```

### **Arguments:**

- rntp*     A pointer to a **timespec** structure that specifies the time interval for which you want to suspend the thread.
- rmtp*     NULL, or a pointer to a **timespec** structure where the function can store the amount of time remaining in the interval (the requested time minus the time actually slept).

### **Library:**

**libc**

### **Description:**

The *nanosleep()* function causes the calling thread to be suspended from execution until either:

- The time interval specified by the *rntp* argument has elapsed
- Or
- A signal is delivered to the thread, and the signal's action is to invoke a signal-catching function or terminate the process.

The suspension time may be longer than requested because the argument value is rounded up to be a multiple of the system timer resolution or because of scheduling and other system activity.

### **Returns:**

- 0     The requested time has elapsed.
- 1    The *nanosleep()* function was interrupted by a signal (*errno* is set).

**Errors:**

|        |                                                                                                                                                                                               |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | All timers are in use. You'll have to wait for a process to release one.                                                                                                                      |
| EFAULT | A fault occurred trying to access the buffers provided.                                                                                                                                       |
| EINTR  | The <i>nanosleep()</i> function was interrupted by a signal.                                                                                                                                  |
| EINVAL | The number of nanoseconds specified by the <i>tv_nsec</i> member of the <code>timespec</code> structure pointed to by <i>rqtp</i> is less than zero or greater than or equal to 1000 million. |

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno*, *clock\_getres()*, *clock\_gettime()*, *clock\_settime()*, *sleep()*,  
*timer\_create()*, *timer\_delete()*, *timer\_gettime()*, *timer\_settime()*,  
`timespec`

## ***nanospin()***

© 2004, QNX Software Systems Ltd.

*Busy-wait without thread blocking for a period of time*

### **Synopsis:**

```
#include <time.h>

int nanospin(const struct timespec *when);
```

### **Arguments:**

*when*     A pointer to a **timespec** structure that specifies the amount of time to busy-wait for. This is a duration, not an absolute time.

### **Library:**

**libc**

### **Description:**

The *nanospin()* function occupies the CPU for the amount of time specified by the argument *when* without blocking the calling thread. (The thread isn't taken off the ready list.) The function is essentially a **do...while** loop.

The *nanospin\*()* functions are designed for use with hardware that requires short time delays between accesses. You should use them only to delay for times less than a few milliseconds.

The first time *nanospin()* is called, the C library invokes *nanospin\_calibrate()*, if you haven't already called it.

### **Returns:**

|        |                                                                                                                      |
|--------|----------------------------------------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                                             |
| E2BIG  | The delay specified by <i>when</i> is greater than 500 milliseconds.                                                 |
| ENOSYS | This system's <b>startup-*</b> program didn't initialize the timing information necessary to use <i>nanospin()</i> . |

## Classification:

QNX Neutrino

### Safety

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | Read the <i>Caveats</i> |
| Signal handler     | Yes                     |
| Thread             | Yes                     |

## Caveats:

Busy-waiting should be used only when absolutely necessary for accessing hardware.

It isn't safe to call this function in an interrupt handler if *nanospin\_calibrate()* hasn't been called yet.

## See also:

*nanosleep()*, *nanospin\_calibrate()*, *nanospin\_count()*, *nanospin\_ns()*, *nanospin\_ns\_to\_count()*, *sched\_yield()*, *sleep()*, **timespec**

## ***nanospin\_calibrate()***

© 2004, QNX Software Systems Ltd.

*Calibrate before calling nanospin\*()*

### **Synopsis:**

```
#include <time.h>

int nanospin_calibrate(int disable);
```

### **Arguments:**

*disable*     1 to disable interrupts during the call to *nanospin\_calibrate()*, or 0 to enable them; see below.

### **Library:**

libc

### **Description:**

The *nanospin\_calibrate()* function performs the calibration for the *nanospin()\** family of delay functions. The first time that you call *nanospin()*, *nanospin\_ns()*, or *nanospin\_ns\_to\_count()*, the C library invokes *nanospin\_calibrate()*, unless you call it directly first.



---

If you don't directly invoke *nanospin\_calibrate()*, the first *nanospin\*()* call in a process will have an overly long delay.

---

Interrupts occurring during *nanospin\_calibrate()* can throw off its timings. If *disable* is 0 (zero), you can prevent this situation by:

- 1     Letting the thread acquire I/O privilege.
- 2     Disabling the interrupts around the *nanospin\_calibrate()* call.

If *disable* is 1 (one), the code disables interrupts around the calibration loop(s). The calling thread is still responsible for obtaining I/O privilege before calling *nanospin\_calibrate()*.



**Returns:**

|       |                                                                        |
|-------|------------------------------------------------------------------------|
| EOK   | Success.                                                               |
| EINTR | A too-high rate of interrupts occurred during the calibration routine. |
| EPERM | The process doesn't have superuser capabilities.                       |

**Examples:**

Busy-wait for 100 nanoseconds:

```
#include <time.h>
#include <sys/syspage.h>

int disable = 0;
unsigned long time = 100;

...
/* Wake up the hardware, then wait for it to be ready. */

if ((nanospin_calibrate(disable)) == EOK)
 nanospin_count(nanospin_ns_to_count(time));
else
 printf ("Didn't calibrate successfully.\n");

/* Use the hardware. */
...
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*nanospin()*, *nanospin\_count()*, *nanospin\_ns()*, *nanospin\_ns\_to\_count()*

**Synopsis:**

```
#include <time.h>

void nanospin_count(unsigned long count);
```

**Arguments:**

*count*     The number of iterations that you want to busy-wait for.

**Library:**

`libc`

**Description:**

The *nanospin\_count()* function busy-waits for the number of iterations specified in *count*. Use *nanospin\_ns\_to\_count()* to turn a number of nanoseconds into an iteration count suitable for *nanospin\_count()*.

The *nanospin\*()* functions are designed for use with hardware that requires short time delays between accesses. You should use them only to delay for times less than a few milliseconds.

**Examples:**

Busy-wait for at least 100 nanoseconds:

```
#include <time.h>
#include <sys/syspage.h>

unsigned long time = 100;

...
/* Wake up the hardware, then wait for it to be ready. */

nanospin_count(nanospin_ns_to_count(time));

/* Use the hardware. */
...
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

You should use busy-waiting only when absolutely necessary for accessing hardware.

**See also:**

*nanosleep(), nanospin(), nanospin\_calibrate(), nanospin\_ns(), nanospin\_ns\_to\_count(), sched\_yield(), sleep()*

## Synopsis:

```
#include <time.h>

int nanospin_ns(unsigned long nsec);
```

## Arguments:

*nsec*      The number of nanoseconds that you want to busy-wait for.

## Library:

`libc`

## Description:

The *nanospin\_ns()* function busy-waits for the number of nanoseconds specified in *nsec*, without blocking the calling thread.

The *nanospin\*()* functions are designed for use with hardware that requires short time delays between accesses. You should use them only to delay for times less than a few milliseconds.

The first time you call *nanospin\_ns()*, the C library invokes *nanospin\_calibrate()*, if you haven't invoked it directly first.

## Returns:

|        |                                                                                                                         |
|--------|-------------------------------------------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                                                |
| E2BIG  | The delay specified by <i>nsec</i> is greater than 500 milliseconds.                                                    |
| ENOSYS | This system's <b>startup-*</b> program didn't initialize the timing information necessary to use <i>nanospin_ns()</i> . |

## Classification:

QNX Neutrino

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | Read the <i>Caveats</i> |
| Signal handler     | Yes                     |
| Thread             | Yes                     |

**Caveats:**

You should use busy-waiting only when absolutely necessary for accessing hardware.

It isn't safe to call this function in an interrupt handler if *nanospin\_calibrate()* hasn't been called yet.

**See also:**

*nanosleep()*, *nanospin()*, *nanospin\_calibrate()*, *nanospin\_count()*, *nanospin\_ns\_to\_count()*, *sched\_yield()*, *sleep()*

## ***nanospin\_ns\_to\_count()***

*Convert a time in nanoseconds into a number of iterations*

### **Synopsis:**

```
#include <time.h>

unsigned long nanospin_ns_to_count(
 unsigned long nsec);
```

### **Arguments:**

*nsec*     The number of nanoseconds that you want to convert.

### **Library:**

`libc`

### **Description:**

The *nanospin\_ns\_to\_count()* function converts the number of nanoseconds specified in *nsec* into an iteration count suitable for *nanospin\_count()*.

The *nanospin\*()* functions are designed for use with hardware that requires short time delays between accesses. You should use them only to delay for times less than a few milliseconds.

The first time that you call *nanospin\_ns\_to\_count()*, the C library invokes *nanospin\_calibrate()* if you haven't invoked it directly first.

### **Returns:**

The amount of time to busy-wait, or -1 if an error occurred (*errno* is set).

### **Errors:**

ENOSYS     This system's **startup-\*** program didn't initialize the timing information necessary to use *nanospin\_ns\_to\_count()*.

## Examples:

Busy-wait for at least one nanosecond:

```
#include <time.h>
#include <sys/syspage.h>

unsigned long time = 1;

...
/* Wake up the hardware, then wait for it to be ready. */

/*
 The C library invokes nanospin_calibrate
 if it hasn't already been called.
*/

nanospin_count(nanospin_ns_to_count(time));

/* Use the hardware. */
...
```

## Classification:

QNX Neutrino

### Safety

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | Read the <i>Caveats</i> |
| Signal handler     | Yes                     |
| Thread             | Yes                     |

## Caveats:

You should use busy-waiting only when absolutely necessary for accessing hardware.

It isn't safe to call this function in an interrupt handler if *nanospin\_calibrate()* hasn't been called yet.



**See also:**

*nanospin(), nanospin\_calibrate(), nanospin\_count(), nanospin\_ns()*

## ***nap()***

© 2004, QNX Software Systems Ltd.

*Sleep for a given number of milliseconds*

### **Synopsis:**

```
#include <unix.h>

unsigned int nap(unsigned int ms);
```

### **Arguments:**

*ms*     The number of milliseconds that you want the process to sleep.

### **Library:**

`libc`

### **Description:**

The *nap()* routine delays the calling process for *ms* milliseconds. This function is the same as *delay()* and is similar to *napms()*.

### **Classification:**

Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **See also:**

*delay()*, *napms()*

**Synopsis:**

```
#include < curses.h>

int napms(int ms);
```

**Arguments:**

*ms*      The number of milliseconds that you want the process to sleep.

**Library:**

`libc`

**Description:**

The *napms()* routine delays the calling process for *ms* milliseconds. This function is similar to *delay()* and *nap()*.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*delay()*, *nap()*

## ***nbaconnect()***

© 2004, QNX Software Systems Ltd.

*Initiate a connection on a socket (nonblocking)*

### **Synopsis:**

```
#include <sys/socket.h>

int nbaconnect(int s,
 const struct sockaddr * name,
 size_t namelen);
```

### **Arguments:**

*s*            The descriptor of the socket on which to initiate the connection.

*name*        The name of the socket to connect to for a SOCK\_STREAM connection.

*namelen*     The length of the *name*, in bytes.

### **Library:**

**libsocket**

### **Description:**

The *nbaconnect()* function is called in place of *connect()*, to prevent a nonblocking *connect()* from blocking during an autoconnect (see */etc/autoconnect*).

When the autoconnect behavior is used, *connect()* may block your application while waiting for the autoconnect to complete; *nbaconnect()* allows your application to continue executing during the autoconnect.

The *nbaconnect()* function takes the same arguments as *connect()*, but it differs in the return value when an autoconnect is required. If an autoconnect is required, a file descriptor (*fd*) is returned. The *fd* is used in the call to *nbaconnect\_result()* to get the *errno* related to the autoconnect and the connect attempt.




---

Since *nbaconnect\_result()* is a blocking call, it's recommended that you call *select()* first to determine if there's data available on the pipe.

---

When the data's available, call *nbaconnect\_result()* to get the status of the *nbaconnect()* attempt.

If an autoconnect isn't required, *nbaconnect()* returns -1 and exhibits the same behavior as *connect()* on nonblocking sockets (e.g. if -1 is returned and *errno* is set to EINPROGRESS, it's possible to do a *select()* for completion by selecting the socket for writing).

## Returns:

A file descriptor that you can pass to *nbaconnect\_result()* to get the result of the *nbaconnect()* attempt, or -1 if an error occurred (*errno* is set).

## Errors:

Any value from the Errors section in *connect()*, as well as:

EINVAL      The socket file descriptor being passed isn't nonblocking.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

The `pipe` manager must be available.

**See also:**

*accept(), bind(), connect(), errno, fcntl(), getsockname(),  
nbaconnect\_result(), open(), pipe(), read(), select(), socket(), write()  
/etc/autoconnect, pipe* in the *Utilities Reference*.

## ***nbaconnect\_result()***

*Get the status of the previous call to nbaconnect()*

### **Synopsis:**

```
#include <sys/socket.h>

int nbaconnect_result(int fd,
 int * error);
```

### **Arguments:**

*fd*        The file descriptor returned by *nbaconnect()*.

*error*     A pointer to a location where the function can store the status.

### **Library:**

**libsocket**

### **Description:**

The *nbaconnect\_result()* function gets the status of the previous *nbaconnect()* call when an *fd* was returned. Since *nbaconnect\_result()* is a blocking call, it's best to test the status of the *fd* with a call to *select()* to verify that the file descriptor is ready to be read.

When there's data available, the status is put in *error*, which may be any of the *errno* values set by *connect()* during an attempt to make a non-blocking connection.

The *fd* is always closed by this function whether or not there's a status to report.

### **Returns:**

0        The call was successful; *error* contains the status.

-1      An error occurred while obtaining the status.

**Errors:**

Any value from the Errors section in *connect()*, as well as:

EBADF        Invalid *fd*.

ENOMSG      There's no data, or not enough data, from the *fd*.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    Yes

Interrupt handler     No

Signal handler        Yes

Thread                 Yes

**See also:**

*connect()*, *nbaconnect()*, *select()*

**autoconnect** in the *Utilities Reference*.



## Synopsis:

```
#include <sys/netmgr.h>

#define ND_NODE_CMP(a,b) ...
```

## Arguments:

*a, b* The node descriptors that you want to compare. You can use either the value 0 or ND\_LOCAL\_NODE to refer to the local node.

## Library:

libc

## Description:

The *ND\_NODE\_CMP()* macro compares two node descriptors.

## Returns:

< 0 The node descriptor *a* is less than *b*.  
0 The descriptors refer to the same machine.  
> 0 The node descriptor *a* is greater than *b*.

## Examples:

```
#include <sys/neutrino.h>

uint32_t nd1, nd2;

if (ND_NODE_CMP(nd1, nd2) == 0) {
 /* Same node */
 ...
} else {
 /* Different nodes */
 ...
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*netmgr\_ndtostr()*, *netmgr\_remote\_nd()*, *netmgr\_strtond()*

Qnet Networking chapter of the *Programmer's Guide*

Qnet Networking chapter of the *System Architecture* guide

**Synopsis:**

```
#include <netdb.h>

struct netent {
 char * n_name;
 char ** n_aliases;
 int n_addrtype;
 uint32_t n_net;
};
```

**Description:**

This structure holds information from the network database, */etc/networks*.

The members of this structure are:

|                   |                                                                         |
|-------------------|-------------------------------------------------------------------------|
| <i>n_name</i>     | The name of the network.                                                |
| <i>n_aliases</i>  | A zero-terminated list of alternate names for the network.              |
| <i>n_addrtype</i> | The type of the network number returned; currently only AF_INET.        |
| <i>n_net</i>      | The network number. Network numbers are returned in machine-byte order. |

**Classification:**

Unix, POSIX 1003.1-2001

**See also:**

*endnetent()*, *getnetbyaddr()*, *getnetbyname()*, *getnetent()*, *setnetent()*  
*/etc/networks* in the *Utilities Reference*

## ***netmgr\_ndtostr()***

© 2004, QNX Software Systems Ltd.

*Convert a node descriptor into a string*

### **Synopsis:**

```
#include <sys/netmgr.h>

int netmgr_ndtostr(unsigned flags,
 int nd,
 char * buf,
 size_t maxbuf);
```

### **Arguments:**

- flags* Which part(s) of the Fully Qualified Path Name (FQPN) to adjust; see below.
- nd* The node descriptor that you want to convert.
- buf* A pointer to a buffer where the function can store the converted identifier.
- maxbuf* The size of the buffer.

### **Library:**


`libc`

### **Description:**

The *netmgr\_ndtostr()* function converts a node descriptor, *nd*, to a string and stores the string in the buffer pointed to by *buf*. The size of the buffer is given by *maxbuf*.

A node descriptor is a temporary numeric description of a remote node. For more information, see the Qnet Networking chapter of the *System Architecture* guide.

The *flags* argument indicates which part(s) of the Fully Qualified Path Name (FQPN) to adjust. The following diagram shows the components for the FQPN, `/net/qnet.qnx.com~preferred`:




---

*Components of a fully qualified pathname.*

The Fully Qualified Node Name (FQNN) is `qnet.qnx.com`.

The default string that `netmgr_ndtostr()` builds is the FQNN, plus the Quality of Service (QoS) if it isn't the default (`~loadbalance`). You can pass this string to any other node that can call `netmgr_strtnd()` and has a `nd` that refers to the same node with the same QoS.

A bitwise OR of *flags* modify the default string in the following way:

`ND2S_DIR_HIDE`

Never include the directory.

`ND2S_DIR_SHOW`

Always build a pathname to the root directory (i.e. `/`) of the node indicated by `nd`. For example, calling:

```
netmgr_ndtostr(ND2S_DIR_SHOW, nd, buf, sizeofbuf)
```

on a node with a default domain of `qnx.com` using a `nd` that refers to a FQNN of `peterv.qnx.com` results in the string:

```
/net/peterv.qnx.com/
```

If Qnet isn't active on the node, and `netmgr_ndtostr(ND2S_DIR_SHOW, nd, buf, sizeofbuf)` has a `nd` of `ND_LOCAL_NODE`, then the string is `/`.

ND2S\_DOMAIN\_HIDE

Never include the domain.

ND2S\_DOMAIN\_SHOW

Always include the domain.

ND2S\_LOCAL\_STR

Display shorter node names on your local node. For example, calling:

```
netmgr_ndtostr(ND2S_LOCAL_STR, nd, buf, sizeofbuf)
```

on a node with a default domain of **qnx.com** using a *nd* that refers to an FQPN of **/net/peterv.qnx.com** results in a string of:

```
peterv
```

Whereas, a *nd* that refers to **/net/peterv.anotherhost.com** results in:

```
peterv.anotherhost.com
```

ND2S\_NAME\_HIDE

Never include the name.

ND2S\_NAME\_SHOW

Always include the name.

ND2S\_QOS\_HIDE

Never include the quality of service (QoS).

ND2S\_QOS\_SHOW

Always include the QoS.

**ND2S\_SEP\_FORCE**

Always include a leading separator. For example, calling:

```
netmgr_ndtostr (ND2S_SEP_FORCE|ND2S_DIR_HIDE|ND2S_NAME_HIDE|ND2S_DOMAIN_HIDE|ND2S_QOS_SHOW, nd, buf, sizeofbuf)
```

with a *nd* of `ND_LOCAL_NODE` results in a string of:

```
~loadbalance
```

This is useful if you want to concatenate each component of the FQPN individually.



---

Don't use a `ND2S*_HIDE` and a corresponding `ND2S*_SHOW` together.

---

**Returns:**

The length of the string, or -1 if an error occurs (*errno* is set).

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/netmgr.h>

int main ()
{
 int nd1, nd2, nd3, len;
 char path1[50] = "/net/dave",
 path2[50] = "/net/karen",
 buff[100];

 nd1 = netmgr_strttond(path1, NULL);
 if (nd1 == -1) {
 perror ("netmgr_strttond");
 return EXIT_FAILURE;
 }
 else {
 printf ("Node id for %s is %d.\n", path1, nd1);
 }
}
```

```
nd2 = netmgr_strtond(path2, NULL);
if (nd2 == -1) {
 perror ("netmgr_strtond");
 return EXIT_FAILURE;
}
else {
 printf ("Node id for %s is %d.\n", path2, nd2);
}

nd3 = netmgr_remote_nd (nd2, nd1);
if (nd3 == -1) {
 perror ("netmgr_strtond");
 return EXIT_FAILURE;
}
else {
 printf ("Node id for %s, relative to %s, is %d.\n",
 path1, path2, nd3);
}

len = netmgr_ndtostr (ND2S_DIR_HIDE | ND2S_DOMAIN_SHOW |
 ND2S_NAME_SHOW | ND2S_QOS_SHOW, nd1, buff, 100);
if (len == -1) {
 perror ("netmgr_ndtostr");
}
else {
 printf ("Node name for %d is %s.\n", nd1, buff);
 return EXIT_FAILURE;
}

return EXIT_SUCCESS;
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*ND\_NODE\_CMP()*, *netmgr\_remote\_nd()*, *netmgr\_strtond()*

Qnet Networking chapter of the *Programmer's Guide*

Qnet Networking chapter of the *System Architecture* guide

## ***netmgr\_remote\_nd()***

© 2004, QNX Software Systems Ltd.

*Get a node descriptor that's relative to a remote node*

### **Synopsis:**

```
#include <sys/netmgr.h>

int netmgr_remote_nd(int remote_nd,
 int local_nd);
```

### **Arguments:**

*remote\_nd*     The node descriptor of a remote node.

*local\_nd*     A node descriptor, relative to the local node, that you want to convert to be relative to the remote node.

### **Library:**

`libc`

### **Description:**

The *netmgr\_remote\_nd()* function converts a node descriptor that's relative to the local node into a node descriptor that's relative to the specified remote node.

### **Returns:**

The node descriptor, relative to the remote node, or -1 if an error occurred (*errno* is set).

### **Examples:**

See *netmgr\_ndtostr()*.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ND\_NODE\_CMP()*, *netmgr\_ndtostr()*, *netmgr\_strtond()*

Qnet Networking chapter of the *Programmer's Guide*

Qnet Networking chapter of the *System Architecture* guide

## ***netmgr\_strtond()***

© 2004, QNX Software Systems Ltd.

*Convert a string into a node descriptor*

### **Synopsis:**

```
#include <sys/netmgr.h>

int netmgr_strtond(const char * nodename,
 char ** endstr);
```

### **Arguments:**

|                 |                                                                                                                               |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------|
| <i>nodename</i> | The string that you want to convert into a node descriptor.                                                                   |
| <i>endstr</i>   | NULL, or the address of a location where the function can store a pointer to the character after the node name in the string. |

### **Library:**

libc

### **Description:**

The *netmgr\_strtond()* function converts a string to a node descriptor. If *endstr* isn't NULL, it's set to point to the character after the node name in the given string.

### **Returns:**

The node descriptor, or -1 if an error occurred (*errno* is set).

### **Errors:**

ENOTSUP      Qnet isn't running.

### **Examples:**

See *netmgr\_ndtostr()*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*ND\_NODE\_CMP()*, *netmgr\_ndtostr()*, *netmgr\_remote\_nd()*

Qnet Networking chapter of the *Programmer's Guide*

Qnet Networking chapter of the *System Architecture* guide

## ***nextafter()*, *nextafterf()***

© 2004, QNX Software Systems Ltd.

*Compute the next representable double-precision floating-point number*

### **Synopsis:**

```
#include <math.h>

double nextafter (double x,
 double y);

float nextafterf (float x,
 float y);
```

### **Arguments:**

- x*     The number that you want the next number after.
- y*     A number that specifies the direction you want to go; see below.

### **Library:**

`libm`

### **Description:**

The *nextafter()* and *nextafterf()* functions compute the next representable double-precision floating-point value following *x* in the direction of *y*.

### **Returns:**

The next machine floating-point number of *x* in the direction towards *y*.

| <b>If:</b> | <b><i>nextafter()</i> returns:</b> |
|------------|------------------------------------|
|------------|------------------------------------|

|                 |                                                              |
|-----------------|--------------------------------------------------------------|
| $y < x$         | The next possible floating-point value less than <i>y</i>    |
| $y > x$         | The next possible floating-point value greater than <i>x</i> |
| <i>x</i> is NAN | NAN                                                          |

*continued...*

**If:** *nextafter()* returns:

---

y is NAN    NAN

x is finite     $\pm$ HUGE\_VAL, according to the sign of x (*errno* is set)

---



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Examples:**

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

void dump_to_hex(double d) {
 printf("0x%08x %08x \n",
 (uint32_t)((uint64_t*)&d >> 32),
 (uint32_t)((uint64_t*)&d));
}

int main(int argc, char** argv)
{
 double a, b, c;

 a = 0 ;
 b = nextafter(a, -1);
 c = nextafter(a, 1);
 printf("Next possible value before %f is %f \n", a, b);
 printf("-->"); dump_to_hex(a);
 printf("-->"); dump_to_hex(b);
 printf("Next possible value after %f is %f \n", a, c);
 printf("-->"); dump_to_hex(a);
 printf("-->"); dump_to_hex(c);

 return(0);
}
```

produces the output:

```
Next possible value before 0.000000 is 0.000000
```

```
-->0x00000000 00000000
-->0x80000000 00000001
Next possible value after 0.000000 is 0.000000
-->0x00000000 00000000
```

## Classification:

*nextafter()* is standard Unix; *nextafterf()* is ANSI (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



**Synopsis:**

```
#include <ftw.h>

int nftw(const char *path,
 int (*fn)(const char *fname,
 const struct stat *sbuf,
 int flags,
 struct FTW *ftw) ,
 int depth,
 int flags);
```

**Arguments:**

*path*      The path of the directory whose file tree you want to walk.

*fn*        A pointer to a function that you want to call for each file; see below.

*depth*     The maximum number of file descriptors that *nftw()* can use. The *nftw()* function uses one file descriptor for each level in the tree.

If *depth* is zero or negative, the effect is the same as if it were 1. The *depth* must not be greater than the number of file descriptors currently available for use. The *nftw()* function is faster if *depth* is at least as large as the number of levels in the tree.

*flags*     The value of *flags* is constructed by the bitwise ORing of values from the following list, defined in the `<ftw.h>` header file.

FTW\_CHDIR    If set, *nftw()* changes the current working directory to each directory as it reports files in that directory.

FTW\_DEPTH    If set, *nftw()* reports all files in a directory before reporting the directory itself (otherwise the directory is reported before any file it contains).

- FTW\_MOUNT If set, *nftw()* only reports files on the same filesystem as *path*.
- FTW\_PHYS If set, *nftw()* performs a physical walk and doesn't follow any symbolic link.

## Library:

`libc`

## Description:

The *nftw()* function recursively descends the directory hierarchy identified by *path*. For each object in the hierarchy, *nftw()* calls the user-defined function *fn()*, passing to it:

- a pointer to a NULL-terminated character string containing the name of the object
- a pointer to a `stat` structure (see *stat()*) containing information about the object
- an integer. Possible values of the integer, defined in the `<nftw.h>` header, are:

- FTW\_F The object is a file.
- FTW\_D The object is a directory.
- FTW\_DNR The object is a directory that can't be read. Descendents of the directory aren't processed.
- FTW\_DP The object is a directory, and its contents have been reported. See the FTW\_DEPTH flag above.
- FTW\_NS The *stat()* failed on the object because the permissions weren't appropriate. The `stat` buffer passed to *fn()* is undefined.
- FTW\_SL The object is a symbolic link. See the FTW\_PHYS flag above.
- FTW\_SLN The object is a symbolic link that does not name an existing file.

- a pointer to a FTW structure, which contains the following fields:
  - base* The offset of the objects filename in the pathname passed as the first argument to *fn()*.
  - level* The depth relative to the root of the walk (where the root is level 0).
  - quit* A flag that can be set to control the behaviour of *nftw()* within the current directory. If assigned, it may be given the following values:
    - FTW\_SKR Skip the remainder of this directory
    - FTW\_SKD If the object is FTW\_D, then do not enter into this directory.

The tree traversal continues until the tree is exhausted, an invocation of *fn()* returns a nonzero value, or some error is detected within *nftw()* (such as an I/O error). If the tree is exhausted, *nftw()* returns zero. If *fn()* returns a nonzero value, *nftw()* stops its tree traversal and returns whatever value was returned by *fn()*.

When *nftw()* returns, it closes any file descriptors it opened; it doesn't close any file descriptors that may have been opened by *fn()*.

## Returns:

- 0 Success.
- 1 An error (other than EACCESS) occurred (*errno* is set).

## Classification:

Standard Unix, *nftw64()* is for large-file support

### Safety

Cancellation point Yes

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

Because *nftw()* is recursive, it might terminate with a memory fault when applied to very deep file structures.

This function uses *malloc()* to allocate dynamic storage during its operation. If *nftw()* is forcibly terminated, for example if *longjmp()* is executed by *fn()* or an interrupt routine, *nftw()* doesn't have a chance to free that storage, so it remains permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn()* return a nonzero value at its next invocation.

**See also:**

*ftw()*, *longjmp()*, *malloc()*, *stat()*

## Synopsis:

```
#include <unistd.h>

int nice(int incr);
```

## Arguments:

*incr* The amount that you want to add to the process's priority.

## Library:

`libc`

## Description:

The *nice()* function allows a process to change its priority. The invoking process must be in a scheduling class that supports the operation.

The *nice()* function adds the value of *incr* to the nice value of the calling process. A process's nice value is a nonnegative number; a greater positive value results in a lower CPU priority.

A maximum nice value of  $2 * NZERO - 1$  and a minimum nice value of 0 are imposed by the system. NZERO is defined in `<limits.h>` with a default value of 20. If you request a value above or below these limits, the nice value is set to the corresponding limit. A nice value of 40 is treated as 39. Only a process with superuser privileges can lower the nice value.

## Returns:

The new nice value minus NZERO. If an error occurred, -1 is returned, the process's nice value isn't changed, and *errno* is set.

**Errors:**

- EINVAL The *nice()* function was called by a process in a scheduling class other than time-sharing.
- EPERM The *incr* argument was negative or greater than 40, and the effective user ID of the calling process isn't the superuser.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

As -1 is a permissible return value in a successful situation, an application wishing to check for error situations should set *errno* to 0, then call *nice()*, and if it returns -1, check to see if *errno* is nonzero.

**See also:**

*execl()*, *execle()*, *execlp()*, *execlepe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*

**nice** in the *Utilities Reference*

*Generate a pseudo-random nonnegative long integer in a thread-safe manner*

**Synopsis:**

```
#include <stdlib.h>

long nrand48(unsigned short xsubi[3]);
```

**Arguments:**

*xsubi*     An array that comprises the 48 bits of the initial value that you want to use.

**Library:**

```
libc
```

**Description:**

The *nrand48()* function uses a linear congruential algorithm and 48-bit integer arithmetic to generate a nonnegative **long** integer uniformly distributed over the interval  $[0, 2^{31}]$ .

The *xsubi* array should contain the desired initial value; this makes *nrand48()* thread-safe, and lets you start a sequence of random numbers at any known value.

**Returns:**

A pseudo-random **long** integer.

**Classification:**

Standard Unix

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*drand48(), erand48(), jrand48(), lcong48(), lrand48(), mrand48(),  
seed48(), srand48()*



### **Synopsis:**

```
#include <time.h>

void nsec2timespec(struct timespec *timespec_p,
 _uint64 nsec);
```

### **Arguments:**

*timespec\_p*     A pointer to a **timespec** structure where the function can store the converted time.

*nsec*            The number of nanoseconds that you want to convert.

### **Library:**

**libc**

### **Description:**

This function converts the given number of nanoseconds, *nsec*, into seconds and nanoseconds, and stores them in the **timespec** structure pointed to by *timespec\_p*.

### **Classification:**

Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`timespec`, `timespec2nsec()`

*Convert a 32-bit value from network-byte order to host-byte order*

**Synopsis:**

```
#include <arpa/inet.h>

uint32_t ntohl(uint32_t netlong);
```

**Arguments:**

*netlong*     The value that you want to convert.

**Library:**

`libc`

**Description:**

The *ntohl()* function converts a 32-bit value from network-byte order to host-byte order. If a machine's byte order is the same as the network order, this routine is defined as a null macro.

You most often use this routine in conjunction with internet addresses and ports returned by *gethostbyname()* and *getservent()*.

**Returns:**

The value in host-byte order.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*gethostbyname(), getservent(), htonl(), htons(), ntohs()*

**Synopsis:**

```
#include <arpa/inet.h>

uint16_t ntohs(uint16_t netshort);
```

**Arguments:**

*netshort*     The value that you want to convert.

**Library:**

`libc`

**Description:**

The *ntohs()* function converts a 16-bit value from network-byte order to host-byte order. If a machine's byte order is the same as the network order, this routine is defined as a null macro.

You most often use this routine in conjunction with internet addresses and ports returned by *gethostbyname()* and *getservent()*.

**Returns:**

The value in host-byte order.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*gethostbyname(), getservent(), htonl(), htons(), ntohl()*

## Synopsis:

```
#include <stddef.h>

#define offsetof(composite, name) ...
```

## Arguments:

*composite*     A **struct** or **union**.

*name*            The name of an element in *composite*.

## Library:

libc

## Description:

The *offsetof()* macro returns the offset of the element *name* within the **struct** or **union** *composite*.

This provides a portable method to determine the offset.

## Returns:

The offset of *name*.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <stddef.h>

struct new_def
{
 char *first;
 char second[10];
 int third;
};

int main(void)
{
 printf("first:%d second:%d third:%d\n",
 offsetof(struct new_def, first),
```

```
 offsetof(struct new_def, second),
 offsetof(struct new_def, third));
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

This is a macro.



## **Synopsis:**

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>

int open(const char * path,
 int oflag,
 ...);

int open64(const char * path,
 int oflag,
 ...);
```

## **Arguments:**

*path*     The path name of the file that you want to open.

*oflag*    Flags that specify the status and access modes of the file; see below.

If you set `O_CREAT` in *oflag*, you must also specify the following argument:

*mode*     An object of type `mode_t` that specifies the access mode that you want to use for a newly created file. For more information, see “Access permissions” in the documentation for *stat()*, and the description of `O_CREAT`, below.

## **Library:**

`libc`

## **Description:**

The *open()* and *open64()* functions open the file named by *path*, creating an open file description that refers to the file, and a file descriptor that refers to the file description. The file status flags and

the file access modes of the open file description are set according to the value of *oflag*.



---

These functions ignore any advisory locks that you set with *fcntl()*.

---

The open file descriptor created is new, and therefore isn't shared with any other process in the system.

Construct the value of *oflag* by bitwise ORing values from the following list, defined in the `<fcntl.h>` header file. You must specify exactly one of the first three values (file access modes) below in the value of *oflag*:

- O\_RDONLY      Open for reading only.
- O\_RDWR        Open for reading and writing. Opening a FIFO for read-write is unsupported.
- O\_WRONLY      Open for writing only.

You can also specify any combination of the remaining flags in the value of *oflag*:

- O\_APPEND      If set, the file offset is set to the end of the file prior to each write.
- O\_CLOEXEC     Close the file descriptor on execution.
- O\_CREAT        This option requires a third argument, *mode*, which is of type `mode_t`. If the file exists, this flag has no effect, except in combination with O\_EXCL as noted below.  
  
Otherwise, the file is created; the file's user ID is set to the effective user ID of the process; the group ID is set to the effective group ID of the process or the group ID of the file's parent directory (see *chmod()*).

The permission bits, as defined in `<sys/stat.h>`, are set to the value of *mode*, except those bits set in the process's file mode creation mask (see *umask()* for details). Bits set in *mode* other than the file permission bits (i.e. the file type bits) are ignored. The *mode* argument doesn't affect whether the file is opened for reading, for writing, or for both.

- O\_DSYNC** If set, this flag affects subsequent I/O calls; each call to *write()* waits until all data is successfully transferred to the storage device such that it's readable on any subsequent open of the file (even one that follows a system failure) in the absence of a failure of the physical storage medium. If the physical storage medium implements a non-write-through cache, then a system failure may be interpreted as a failure of the physical storage medium, and data may not be readable even if this flag is set and the *write()* indicates that it succeeded.
- O\_EXCL** If you set both **O\_EXCL** and **O\_CREAT**, *open()* fails if the file exists. The check for the existence of the file and the creation of the file if it doesn't exist are atomic; no other process that's attempting the same operation with the same filename at the same time will succeed. Specifying **O\_EXCL** without **O\_CREAT** has no effect.
- O\_LARGEFILE** Allow the file offset to be 64 bits long.
- O\_NOCTTY** If set, and *path* identifies a terminal device, the *open()* function doesn't cause the terminal device to become the controlling terminal for the process.
- O\_NONBLOCK**
- When opening a FIFO with **O\_RDONLY** or **O\_WRONLY** set:

If `O_NONBLOCK` is set:

Calling *open()* for reading-only returns without delay. Calling *open()* for writing-only returns an error if no process currently has the FIFO open for reading.

If `O_NONBLOCK` is clear:

Calling *open()* for reading-only blocks until a process opens the file for writing. Calling *open()* for writing-only blocks until a process opens the file for reading.

- When opening a block special or character special file that supports nonblocking opens:

If `O_NONBLOCK` is set:

The *open()* function returns without waiting for the device to be ready or available. Subsequent behavior of the device is device-specific.

If `O_NONBLOCK` is clear:

The *open()* function waits until the device is ready or available before returning. The definition of when a device is ready is device-specific.

- Otherwise, the behavior of `O_NONBLOCK` is unspecified.

|                        |                                                                                                                                                                                                                                                                            |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>O_REALIDS</code> | Use the real <code>uid/gid</code> for permissions checking.                                                                                                                                                                                                                |
| <code>O_RSYNC</code>   | Read I/O operations on the file descriptor complete at the same level of integrity as specified by the <code>O_DSYNC</code> and <code>O_SYNC</code> flags.                                                                                                                 |
| <code>O_SYNC</code>    | If set, this flag affects subsequent I/O calls; each call to <i>read()</i> or <i>write()</i> is complete only when both the data has been successfully transferred (either read or written) and all file system information relevant to that I/O operation (including that |

required to retrieve said data) is successfully transferred, including file update and/or access times, and so on. See the discussion of a successful data transfer in O\_DSYNC, above.

**O\_TRUNC** If the file exists and is a regular file, and the file is successfully opened O\_WRONLY or O\_RDWR, the file length is truncated to zero and the mode and owner are left unchanged. O\_TRUNC has no effect on FIFO or block or character special files or directories. Using O\_TRUNC with O\_RDONLY has no effect.

The largest value that can be represented correctly in an object of type `off_t` shall be established as the offset maximum in the open file description.

## Returns:

A nonnegative integer representing the lowest numbered unused file descriptor. On a file capable of seeking, the file offset is set to the beginning of the file. Otherwise, -1 is returned (*errno* is set).




---

In QNX Neutrino, the returned file descriptor is the same as the connection ID (or *coid*) used by the Neutrino-specific functions.

---

## Errors:

**EACCES** Search permission is denied on a component of the *path* prefix, or the file exists and the permissions specified by *oflag* are denied, or the file doesn't exist and write permission is denied for the parent directory of the file to be created.

**EBADFSYS** While attempting to open the named file, either the file itself or a component of the path prefix was found to be corrupted. A system failure — from which no automatic recovery is possible —

occurred while the file was being written to, or while the directory was being updated. You'll need to invoke appropriate systems-administration procedures to correct this situation before proceeding.

|              |                                                                                                                                                                                  |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EBUSY        | File access was denied due to a conflicting open (see <i>sopen()</i> ).                                                                                                          |
| EEXIST       | The O_CREAT and O_EXCL flags are set, and the named file exists.                                                                                                                 |
| EINTR        | The <i>open()</i> operation was interrupted by a signal.                                                                                                                         |
| EINVAL       | The requested synchronized modes (O_SYNC, O_DSYNC, O_RSYNC) aren't supported.                                                                                                    |
| EISDIR       | The named file is a directory, and the <i>oflag</i> argument specifies write-only or read/write access.                                                                          |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                   |
| EMFILE       | Too many file descriptors are currently in use by this process.                                                                                                                  |
| ENAMETOOLONG | The length of the <i>path</i> string exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.                                                                          |
| ENFILE       | Too many files are currently open in the system.                                                                                                                                 |
| ENOENT       | The O_CREAT flag isn't set, and the named file doesn't exist; or O_CREAT is set and either the path prefix doesn't exist, or the <i>path</i> argument points to an empty string. |
| ENOSPC       | The directory or filesystem that would contain the new file can't be extended.                                                                                                   |
| ENOSYS       | The <i>open()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                                                                       |

|           |                                                                                                                                                                                            |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENOTDIR   | A component of the path prefix isn't a directory.                                                                                                                                          |
| ENXIO     | The O_NONBLOCK flag is set, the named file is a FIFO, O_WRONLY is set, no process has the file open for reading, or the media associated with the file has been removed (e.g. CD, floppy). |
| EOVERFLOW | The named file is a regular file and the size of the file can't be represented correctly in an object of type <code>off_t</code> .                                                         |
| EROFS     | The named file resides on a read-only filesystem and either O_WRONLY, O_RDWR, O_CREAT (if the file doesn't exist), or O_TRUNC is set in the <i>oflag</i> argument.                         |

## Examples:

```
#include <sys/types.h>
#include <sys/stat.h>
#include <fcntl.h>
#include <stdlib.h>

int main(void)
{
 int fd;

 /* open a file for output */
 /* replace existing file if it exists */
 /* with read/write perms for owner */
 fd = open("myfile.dat",
 O_WRONLY | O_CREAT | O_TRUNC,
 S_IRUSR | S_IWUSR);

 /* read a file that is assumed to exist */

 fd = open("myfile.dat", O_RDONLY);

 /* append to the end of an existing file */
 /* write a new file if file doesn't exist */
 /* with full read/write permissions */
 fd = open("myfile.dat",
 O_WRONLY | O_CREAT | O_APPEND,
 S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP
```

```
 | S_IROTH | S_IWOTH);
 return EXIT_SUCCESS;
}
```

## Classification:

*open()* is POSIX 1003.1; *open64()* is for large-file support

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The *open()* function includes POSIX 1003.1-1996 and QNX extensions.

## See also:

*chmod()*, *close()*, *creat()*, *dup()*, *dup2()*, *errno*, *fcntl()*, *fstat()*, *lseek()*, *read()*, *write()*



## Synopsis:

```
#include <dirent.h>

DIR * opendir(const char * dirname);
```

## Arguments:

*dirname*      The path of the directory to be opened. It can be relative to the current working directory, or an absolute path.

## Library:

libc

## Description:

The *opendir()* function is used with *readdir()* and *closedir()* to get the list of file names contained in the directory specified by *dirname*.

You can read more than one directory at the same time using the *opendir()*, *readdir()*, *rewinddir()* and *closedir()* functions.



The result of using a directory stream after one of the *exec\*()* or *spawn\*()* functions is undefined. After a call to the *fork()* function, either the parent *or* the child (but not both) can continue processing the directory stream using *readdir()* and *rewinddir()*. If both the parent and child processes use these functions, the result is undefined. Either process can use *closedir()*.

---

## Returns:

A pointer to a **DIR** structure required for subsequent calls to *readdir()* to retrieve the file names in *dirname*, or NULL if *dirname* isn't a valid path (*errno* is set).

**Errors:**

|              |                                                                                                                   |
|--------------|-------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of <i>dirname</i> , or read permission is denied for <i>dirname</i> . |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                    |
| ENAMETOOLONG | The length of <i>dirname</i> exceeds PATH_MAX, or a pathname component is longer than NAME_MAX.                   |
| ENOENT       | The named directory doesn't exist.                                                                                |
| ENOSYS       | The <i>opendir()</i> function isn't implemented for the filesystem specified in <i>dirname</i> .                  |
| ENOTDIR      | A component of <i>dirname</i> isn't a directory.                                                                  |

**Examples:**

Get a list of files contained in the directory `/home/fred`:

```
#include <stdio.h>
#include <stdlib.h>
#include <dirent.h>

int main(void)
{
 DIR* dirp;
 struct dirent* direntp;

 dirp = opendir("/home/fred");
 if(dirp == NULL) {
 perror("can't open /home/fred");
 } else {
 for(;;) {
 direntp = readdir(dirp);
 if(direntp == NULL) break;

 printf("%s\n", direntp->d_name);
 }

 closedir(dirp);
 }

 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*closedir()*, *errno*, *readdir()*, *readdir\_r()*, *rewinddir()*, *seekdir()*,  
*telldir()*

## ***openfd()***

© 2004, QNX Software Systems Ltd.

*Open for private access a file associated with a given descriptor*

### **Synopsis:**

```
#include <unistd.h>

int openfd(int fd,
 int oflag);
```

### **Arguments:**

- fd* A file descriptor associated with the file that you want to open.
- oflag* How you want to open the file; a combination of the following bits:
- O\_RDONLY — permit the file to be only read.
  - O\_WRONLY — permit the file to be only written.
  - O\_RDWR — permit the file to be both read and written.
  - O\_APPEND — cause each record that's written to be written at the end of the file.
  - O\_TRUNC — truncate the file to contain no data.

### **Library:**

libc

### **Description:**

The *openfd()* function opens the file associated with the file descriptor, *fd*. This is similar to *dup()*, except the new *fd* has private access modes and offset. The access mode, *oflag*, must be equal to or more restrictive than the access mode of the source *fd*.

### **Returns:**

A file descriptor, or -1 if an error occurred (*errno* is set).

**Errors:**

|        |                                                                                                                             |
|--------|-----------------------------------------------------------------------------------------------------------------------------|
| EBADF  | Invalid file descriptor <i>fd</i> .                                                                                         |
| EACCES | The access mode specified by <i>oflag</i> isn't equal to or more restrictive than the access mode of the source <i>fd</i> . |
| EBUSY  | Sharing mode ( <i>sflag</i> ) was denied due to a conflicting open (see <i>sopenfd()</i> ).                                 |

**Examples:**

```
#include <unistd.h>
#include <stdlib.h>
#include <fcntl.h>

int main (void)
{
 int fd, fd2, oflag;

 fd = open ("etc/passwd", O_RDONLY);
 fd2 = openfd (fd, O_RDONLY);
 return EXIT_SUCCESS;
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*dup(), sopenfd()*

**Synopsis:**

```
#include <syslog.h>

void openlog(const char * ident,
 int logopt,
 int facility);
```

**Arguments:**

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ident</i>    | A string that you want to prepend to every message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <i>logopt</i>   | A bit field specifying logging options; a combination of one or more of the following values with an OR operation:<br><br>LOG_CONS      If <i>syslog()</i> can't pass the message to <b>syslogd</b> , it attempts to write the message to the <b>/dev/console</b> device. The <b>/dev/console</b> device is usually a symlink (see the <b>ln</b> command) to a real device (e.g. <b>/dev/text</b> , <b>/dev/con1</b> or <b>/dev/ser1</b> ).<br><br>LOG_NDELAY    Open the connection to <b>syslogd</b> immediately. Normally the opening is delayed until the first message is logged.<br><br>LOG_PERROR    Write the message to standard error output as well to the system log.<br><br>LOG_PID       Log the process ID with each message. This is useful for identifying instantiations of daemons. |
| <i>facility</i> | Encode a default facility to be assigned to all messages that don't have an explicit facility encoded. In the following list, parameter values marked with an asterisk (*) aren't used by any of the QNX Neutrino standard utilities.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

|                                 |                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------|
| LOG_AUTH *                      | Authorization system.                                                                                      |
| LOG_AUTHPRIV *                  | Same as LOG_AUTH, but logged to a file readable only by selected individuals.                              |
| LOG_CRON *                      | Clock daemon.                                                                                              |
| LOG_DAEMON                      | System daemons (such as <b>routed</b> ) that aren't explicitly provided for by other facilities.           |
| LOG_FTPD                        | File transfer protocol daemon.                                                                             |
| LOG_KERN *                      | Messages generated by the kernel. These can't be generated by any user processes.                          |
| LOG_LPR                         | Line printer spooling system.                                                                              |
| LOG_MAIL                        | Mail system.                                                                                               |
| LOG_NEWS *                      | Network news system.                                                                                       |
| LOG_SYSLOG                      | Messages generated internally by <b>syslogd</b> .                                                          |
| LOG_USER*                       | Messages generated by random user processes. This is the default facility identifier if none is specified. |
| LOG_UUCP *                      | The uucp system.                                                                                           |
| LOG_LOCAL0 through LOG_LOCAL7 * | Reserved for local use.                                                                                    |

**Library:**

`libc`

**Description:**

The *openlog()* function opens the system log and provides for more specialized processing of the messages sent by *syslog()* and *vsyslog()*.



**Examples:**

See *syslog()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*closelog()*, *setlogmask()*, *syslog()*, *vsyslog()*

**logger**, **syslogd** in the *Utilities Reference*

# ***openpty()***

© 2004, QNX Software Systems Ltd.

*Find an available pseudo-tty*

---

## **Synopsis:**

```
#include <unix.h>

int openpty(int* amaster,
 int* aslave,
 char* name,
 struct termios* termp,
 struct winsize* winp);
```

## **Arguments:**

|                |                                                                                                                                                 |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>amaster</i> | A pointer to a location where <i>forkpty()</i> can store the file descriptor of the master side of the pseudo-tty.                              |
| <i>aslave</i>  | A pointer to a location where <i>forkpty()</i> can store the file descriptor of the slave side of the pseudo-tty.                               |
| <i>name</i>    | NULL, or a pointer to a buffer where <i>forkpty()</i> can store the filename of the slave side of the pseudo-tty.                               |
| <i>termp</i>   | NULL, or a pointer to a <b>termios</b> structure that describes the terminal's control attributes to apply to the slave side of the pseudo-tty. |
| <i>winp</i>    | A pointer to a <b>winsize</b> structure that defines the window size to use for the slave side of the pseudo-tty.                               |

## **Library:**

**libc**

## **Description:**

The *openpty()* function finds and opens an available pseudo-tty.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- ENOENT There are no ttys available.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*fork()*, *forkpty()*, *login\_tty()*, **termios**

## ***out8()***

© 2004, QNX Software Systems Ltd.

*Write an 8-bit value to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void out8(uintptr_t port,
 uint8_t val);
```

### **Arguments:**

*port*     The port you want to write the value to.

*val*      The value that you want to write.

### **Library:**

`libc`

### **Description:**

The *out8()* function writes an 8-bit value, specified by *val*, to the specified *port*.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

## **See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*,  
*out8s()*, *out16()*, *out16s()*, *out32()*, *out32s()*

## ***out8s()***

© 2004, QNX Software Systems Ltd.

*Write 8-bit values to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void * out8s(const void * buff,
 unsigned len,
 uintptr_t port);
```

### **Arguments:**

*val*      A pointer to a buffer that holds the values that you want to write.

*len*      The number of values that you want to write.

*port*     The port you want to write the values to.

### **Library:**

`libc`

### **Description:**

The *out8s()* function writes *len* 8-bit values from the buffer pointed to by *buff* to the specified *port*.

### **Returns:**

A pointer to the end of the written data.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point    No

Interrupt handler      Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*,  
*out8()*, *out16()*, *out16s()*, *out32()*, *out32s()*

## ***out16()*, *outbe16()*, *outle16()***

© 2004, QNX Software Systems Ltd.

*Write a 16-bit value to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void out16(uintptr_t port,
 uint16_t val);

#define outbe16(port,
 val) ...

#define outle16(port,
 val) ...
```

### **Arguments:**

*port*     The port you want to write the value to.

*val*       The value that you want to write.

### **Library:**

libc

### **Description:**

The *out16()* function writes the native-endian 16-bit value, specified by *val*, to the specified *port*.

The *outbe16()* and *outle16()* macros write the native-endian 16-bit value, specified by *val*, to the specified *port* in big-endian or little-endian format, respectively.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point   No

*continued...*



**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO.TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

*outbe16()* and *outle16()* are implemented as macros.

**See also:**

*in8(), in8s(), in16(), in16s(), in32(), in32s(), mmap\_device\_io(), out8(), out8s(), out16s(), out32(), out32s()*

## ***out16s()***

© 2004, QNX Software Systems Ltd.

*Write words to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void * out16s(const void * buff,
 unsigned len,
 uintptr_t port);
```

### **Arguments:**

*val*      A pointer to a buffer that holds the values that you want to write.

*len*      The number of values that you want to write.

*port*     The port you want to write the values to.

### **Library:**

`libc`

### **Description:**

The *out16s()* function writes *len* words from the buffer pointed to by *buff* to the specified *port*.

### **Returns:**

A pointer to the end of the written data.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point    No

Interrupt handler      Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*,  
*out8()*, *out8s()*, *out16()*, *out32()*, *out32s()*

## ***out32()*, *outbe32()*, *outle32()***

© 2004, QNX Software Systems Ltd.

*Write a 32-bit value to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void out32(uintptr_t port,
 uint32_t val);

#define outbe16(port,
 val) ...

#define outle32(port,
 val) ...
```

### **Arguments:**

*port*     The port you want to write the value to.

*val*     The value that you want to write.

### **Library:**

libc

### **Description:**

The *out32()* function writes the 32-bit value, specified by *val*, to the specified *port*.

The *outbe32()* and *outle32()* functions macros the native-endian 32-bit value, specified by *val*, to the specified *port* in big-endian or little-endian format, respectively.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point   No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO.TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

*outbe16()* and *outle16()* are implemented as macros.

**See also:**

*in8(), in8s(), in16(), in16s(), in32(), in32s(), mmap\_device\_io(), out8(), out8s(), out16(), out16s(), out32s()*

## ***out32s()***

© 2004, QNX Software Systems Ltd.

*Write longs to a port*

### **Synopsis:**

```
#include <hw/inout.h>

void * out32s(const void * buff,
 unsigned len,
 uintptr_t port);
```

### **Arguments:**

*val*      A pointer to a buffer that holds the values that you want to write.

*len*      The number of values that you want to write.

*port*     The port you want to write the values to.

### **Library:**

`libc`

### **Description:**

The *out32s()* function writes *len* longs from the buffer pointed to by *buff* to the specified *port*.

### **Returns:**

A pointer to the end of the written data.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point    No

Interrupt handler      Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

The calling thread must have I/O privileges; see *ThreadCtl()*'s `_NTO_TCTL_IO` command for details.

The calling process must also use *mmap\_device\_io()* to access the device's I/O registers.

**See also:**

*in8()*, *in8s()*, *in16()*, *in16s()*, *in32()*, *in32s()*, *mmap\_device\_io()*,  
*out8()*, *out8s()*, *out16()*, *out16s()*, *out32()*

## ***pathconf()***

© 2004, QNX Software Systems Ltd.

*Return the value of a configurable limit*

### **Synopsis:**

```
#include <unistd.h>

long pathconf(const char* path,
 int name);
```

### **Arguments:**

*path*      The name of the file whose limit you want to get.

*name*      The name of the configurable limit; see below.

### **Library:**

`libc`

### **Description:**

The *pathconf()* function returns a value of a configurable limit indicated by *name*, which is associated with the filename given in *path*.

Configurable limits are defined in `<confname.h>`, and include at least the following values:

`_PC_LINK_MAX`

Maximum value of a file's link count.

`_PC_MAX_CANON`

Maximum number of bytes in a terminal's canonical input buffer (edit buffer).

`_PC_MAX_INPUT`

Maximum number of bytes in a terminal's raw input buffer.

`_PC_NAME_MAX`

Maximum number of bytes in a file name (not including the terminating null).



**\_PC\_PATH\_MAX**

Maximum number of bytes in a pathname (not including the terminating null).

**\_PC\_PIPE\_BUF**

Maximum number of bytes that can be written atomically when writing to a pipe.

**\_PC\_CHOWN\_RESTRICTED**

If defined (not -1), indicates that the use of the *chown()* function is restricted to a process with **root** privileges, and to changing the group ID of a file to the effective group ID of the process or to one of its supplementary group IDs.

**\_PC\_NO\_TRUNC**

If defined (not -1), indicates that the use of pathname components longer than the value given by **\_PC\_NAME\_MAX** will generate an error.

**\_PC\_VDISABLE**

If defined (not -1), this is the character value which can be used to individually disable special control characters in the **termios** control structure.

**Returns:**

The requested configurable limit, or -1 if an error occurs (*errno* is set).

**Errors:**

- EACCES** Search permission is denied for a component of *path*.
- EINVAL** The *name* argument is invalid, or the indicated limit isn't supported.
- ELOOP** Too many levels of symbolic links or prefixes.
- ENAMETOOLONG**  
The *path* argument, or a component of *path*, is too long.

- ENOENT      The file doesn't exist.
- ENOSYS      The *pathconf()* function isn't implemented for the filesystem specified in *path*.
- ENOTDIR     A component of the path prefix isn't a directory.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(void)
{
 long value;

 value = pathconf("/dev/con1", _PC_MAX_INPUT);
 printf("Input buffer size is %ld bytes\n",
 value);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*chown()*, *confstr()*, *errno*, *fpathconf()*, *sysconf()*, **termios**  
**getconf** in the *Utilities Reference*

Understanding System Limits chapter of the Neutrino *User's Guide*

## ***pathfind()*, *pathfind\_r()***

© 2004, QNX Software Systems Ltd.

*Search for a file in a list of directories*

### **Synopsis:**

```
#include <libgen.h>

char *pathfind(const char *path,
 const char *name,
 const char *mode);

char *pathfind_r(const char *path,
 const char *name,
 const char *mode,
 char *buff,
 size_t buff_size);
```

### **Arguments:**

|             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-----------|----------|-----------|----------|-------------|----------|--------------|----------|----------------|----------|--------------------|----------|------------|----------|--------------|----------|------------------|----------|-------------------|
| <i>path</i> | A string that specifies the list of the directories that you want to search. The directories named in <i>path</i> are separated by colons.                                                                                                                                                                                                                                                                                                                                                                                                                   |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <i>name</i> | The name of the file you're looking for. If <i>name</i> begins with a slash, the name is treated as an absolute pathname, and <i>path</i> is ignored.                                                                                                                                                                                                                                                                                                                                                                                                        |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <i>mode</i> | A string of option letters chosen from:<br><br><table><tr><td><b>r</b></td><td>Readable.</td></tr><tr><td><b>w</b></td><td>Writable.</td></tr><tr><td><b>x</b></td><td>Executable.</td></tr><tr><td><b>f</b></td><td>Normal file.</td></tr><tr><td><b>b</b></td><td>Block special.</td></tr><tr><td><b>c</b></td><td>Character special.</td></tr><tr><td><b>d</b></td><td>Directory.</td></tr><tr><td><b>p</b></td><td>FIFO (pipe).</td></tr><tr><td><b>u</b></td><td>Set user ID bit.</td></tr><tr><td><b>g</b></td><td>Set group ID bit.</td></tr></table> | <b>r</b> | Readable. | <b>w</b> | Writable. | <b>x</b> | Executable. | <b>f</b> | Normal file. | <b>b</b> | Block special. | <b>c</b> | Character special. | <b>d</b> | Directory. | <b>p</b> | FIFO (pipe). | <b>u</b> | Set user ID bit. | <b>g</b> | Set group ID bit. |
| <b>r</b>    | Readable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>w</b>    | Writable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>x</b>    | Executable.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>f</b>    | Normal file.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>b</b>    | Block special.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>c</b>    | Character special.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>d</b>    | Directory.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>p</b>    | FIFO (pipe).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>u</b>    | Set user ID bit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |
| <b>g</b>    | Set group ID bit.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |          |           |          |           |          |             |          |              |          |                |          |                    |          |            |          |              |          |                  |          |                   |

- k** Sticky bit.
- s** Size nonzero.

*buff* (*pathfind\_r()* only) A pointer to a buffer where *pathfind\_r()* can store the path of the file found.

*buff\_size* (*pathfind\_r()* only) The size of the buffer that *buff* points to.

## Library:

`libc`

## Description:

The *pathfind()* function searches the directories named in *path* for the file *name*. The *pathfind\_r()* function is a thread-safe version of *pathfind()*.

Options read, write, and execute are checked relative to the real (not the effective) user ID and group ID of the current process.

If the file *name*, with all the characteristics specified by *mode*, is found in any of the directories specified by *path*, then these functions return a pointer to a string containing the member of *path*, followed by a slash character (/), followed by *name*.

An empty path member is treated as the current directory. If *name* is found in the current directory, a slash isn't prepended to it; the unadorned name is returned.

The *pathfind\_r()* also includes a buffer, *buff*, and its size, *buff\_size*. This buffer is used to hold the path of the file found.

## Returns:

The path found, or NULL if the file couldn't be found.

## Examples:

Find the `ls` command using the **PATH** environment variable:

```
pathfind (getenv ("PATH"), "ls", "rx");
```

## Classification:

Unix

### *pathfind()*

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

### *pathfind\_r()*

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The string pointed to by the returned pointer is stored in an area that's reused on subsequent calls to *pathfind()*. Don't free this string.

Use *pathfind\_r()* in multithreaded applications.

**See also:**

*access()*, *getenv()*, *mknod()*, *stat()*

**sh** in the *Utilities Reference*

# ***pathmgr\_symlink()***

© 2004, QNX Software Systems Ltd.

*Create a symlink*

## **Synopsis:**

```
#include <sys/pathmgr.h>

int pathmgr_symlink(const char * symlink,
 const char * path);
```

## **Arguments:**

*symlink*     The name of the link that you want to create.

*path*        The path that you want to link to.

## **Library:**

libc

## **Description:**

The *pathmgr\_symlink()* function creates a symbolic link, *path*, in the process manager that redirects to the path specified by *symlink*.

The *pathmgr\_unlink()* function removes the link.



---

The symbolic link isn't permanent and is lost when the system reboots.

---

## **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

## **Examples:**

```
#include <stdio.h>
#include <sys/pathmgr.h>

int main(int argc, char **argv) {

 /* Create a link /mytmp --> /dev/shmem */
```



```
 if(pathmgr_symlink("/dev/shmem", "/mytmp") == -1) {
 perror("Can't make link");
 }

 getchar();
 if(pathmgr_unlink("/mytmp") == -1) {
 perror("Can't unlink ");
 }

 return 0;
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pathmgr\_unlink()*, *symlink()*, *unlink()*

## ***pathmgr\_unlink()***

© 2004, QNX Software Systems Ltd.

*Remove a link*

### **Synopsis:**

```
#include <sys/pathmgr.h>

int pathmgr_unlink(const char * path);
```

### **Arguments:**

*path*     The link that you want to remove.

### **Library:**

`libc`

### **Description:**

The *pathmgr\_unlink()* function removes the link created by *pathmgr\_symlink()*.

### **Returns:**

0        Success.  
-1       An error occurred.

### **Examples:**

See *pathmgr\_symlink()*.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pathmgr\_symlink(), symlink(), unlink()*

## ***pause()***

© 2004, QNX Software Systems Ltd.

*Suspend the process until delivery of a signal*

### **Synopsis:**

```
#include <unistd.h>

int pause(void);
```

### **Library:**

```
libc
```

### **Description:**

The *pause()* function suspends the calling process until delivery of a signal whose action is either to execute a signal handler or to terminate the process.

If the action is to terminate the process, *pause()* doesn't return. If the action is to execute a signal handler, *pause()* returns after the signal handler returns.

### **Returns:**

On error, *pause()* returns -1 and sets *errno* to *EINTR*; otherwise, it never returns.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(void)
{
 /* set an alarm to go off in 5 seconds */
 alarm(5);

 /*
 * Wait until we receive a SIGALRM signal. However,
 * since we don't have a signal handler, any signal
 * will kill us.
 */
 printf("Hang around, "
 " waiting to die in 5 seconds\n");
```

```
 pause();
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*alarm(), errno, sigaction()*

## ***pccard\_arm()***

© 2004, QNX Software Systems Ltd.

*Arm the devp-pccard server*

### **Synopsis:**

```
#include <sys/pccard.h>

int pccard_arm(pccard_t handle,
 int devtype,
 unsigned event,
 int coid);
```

### **Arguments:**

- handle*      The handle returned by *pccard\_attach()*.
- devtype*     The type of device that your application wants to be informed about. Valid devices are:
- \_PCCARD\_DEV\_AIMS — Auto Incrementing Mass Storage.
  - \_PCCARD\_DEV\_ALL — all devices.
  - \_PCCARD\_DEV\_FIXED\_DISK — any hard drive.
  - \_PCCARD\_DEV\_GPIB — General Purpose Interface Bus card.
  - \_PCCARD\_DEV\_MEMORY — memory type device.
  - \_PCCARD\_DEV\_NETWORK — any network adapter.
  - \_PCCARD\_DEV\_PARALLEL — PC parallel device.
  - \_PCCARD\_DEV\_SCSI — any SCSI interface.
  - \_PCCARD\_DEV\_SERIAL — 16450 serial device.
  - \_PCCARD\_DEV\_SOUND — any sound adapter.
  - \_PCCARD\_DEV\_VIDEO — any video adapter.
- event*       The type of event that you want to be notified of:
- **\_PCCARD\_ARM\_INSERT\_REMOVE** — card insertion/removal.
- coid*        A connection ID, obtained from *ConnectAttach()*, that's used to send the pulse.

## Library:

libpccard

## Description:

The *pccard\_arm()* function call is used to request that the **devp-pccard** server notify the user application, via a pulse, when the specified event occurs.

## Returns:

0      Success.  
-1     An error occurred (*errno* is set).

## Errors:

EBADF    Invalid *handle* parameter.

## Examples:

```
/*
 * Ask to be informed when a Network card is inserted
 */

#include <stdio.h>
#include <stdlib.h>
#include <fcntl.h>
#include <sys/neutrino.h>
#include <sys/pccard.h>

int main (void)
{
 pccard_t handle;
 int chid, coid;
 char buf [10];
 struct _pccard_info io;

 if ((handle = pccard_attach (0)) == -1) {
 printf ("Unable to attach to PCCARD\n");
 exit (EXIT_FAILURE);
 }

 if ((chid = ChannelCreate (_NTO_CHF_FIXED_PRIORITY)) == -1) {
```

```
 printf ("Unable to create channel\n");
 exit (EXIT_FAILURE);
}

if ((coid = ConnectAttach (0, 0, chid, _NTO_SIDE_CHANNEL,
 0)) == -1) {
 printf ("Unable to ConnectAttach\n");
 exit (EXIT_FAILURE);
}

if (pccard_arm (handle, _PCCARD_DEV_NETWORK,
 _PCCARD_ARM_INSERT_REMOVE, coid) == -1) {
 perror ("Arm failed");
 exit (EXIT_FAILURE);
}

/* To be informed about any card insertion/removal event,
 * change _PCCARD_DEV_NETWORK to _PCCARD_DEV_ALL.
 */

/*
 * MsgReceive (chid,);
 * Other user logic...
 */

/* Get information from socket 0 - function 0 */
if (pccard_info (handle, 0, &io, sizeof (io)) == -1) {
 perror ("Info failed");
 exit (EXIT_FAILURE);
}

if (io.flags & _PCCARD_FLAG_CARD) {
 printf ("Card inserted in socket 1 - Type %x\n",
 io.window [0].device & 0xff00);
/* Now lock the card in socket 1 with exclusive access */
 if (pccard_lock (handle, 0, 0, O_RDWR | O_EXCL) == -1) {
 perror ("Lock failed");
 exit (EXIT_FAILURE);
 }
/* Read 2 bytes of the CIS from offset 0 in attribute memory */
 if (pccard_raw_read (handle, 0, _PCCARD_MEMTYPE_ATTRIBUTE,
 0, 2, buf) == -1) {
 perror ("Raw read");
 exit (EXIT_FAILURE);
 }
 /* More user logic... */
}
pccard_unlock (handle, 0, 0);
pccard_detach (handle);

return (EXIT_SUCCESS);
```



}

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pccard\_attach()*, *pccard\_detach()*, *pccard\_info()*, *pccard\_lock()*,  
*pccard\_raw\_read()*, *pccard\_unlock()*

## ***pccard\_attach()***

© 2004, QNX Software Systems Ltd.

*Attach to the devp-pccard server*

### **Synopsis:**

```
#include <sys/pccard.h>

pccard_t pccard_attach(int reserved);
```

### **Arguments:**

*reserved*     Pass 0 for this argument.

### **Library:**

`libpccard`

### **Description:**

The *pccard\_attach()* function attaches a user application to the `devp-pccard` server. You must call this function before using any of the other PC card functions, because it returns a handle that all the other PC Card functions use.

### **Returns:**

- >0     A value to be used as *handle* in all other PC Card function calls.
- 1     Can't locate the `devp-pccard` server.
- 2     Send to `devp-pccard` server failed.
- 3     The `devp-pccard` server returned an error (*errno* is set).

### **Errors:**

EBUSY     The `devp-pccard` server is unable to service this request.

**Examples:**

See *pccard\_arm()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_detach()*, *pccard\_info()*, *pccard\_lock()*,  
*pccard\_raw\_read()*, *pccard\_unlock()*

## ***pccard\_detach()***

© 2004, QNX Software Systems Ltd.

*Detach from the devp-pccard server*

### **Synopsis:**

```
#include <sys/pccard.h>

int pccard_detach(pccard_t handle);
```

### **Arguments:**

*handle*     The handle returned by *pccard\_attach()*.

### **Library:**

libpccard

### **Description:**

The *pccard\_detach()* function detaches the user application from the **devp-pccard** server. Any locks that you previously applied with *pccard\_lock()* are freed.

### **Returns:**

0     Success.

-1    An error occurred (*errno* is set).

### **Errors:**

EBADF     Invalid *handle* parameter.

### **Examples:**

See *pccard\_arm()*.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_attach()*, *pccard\_info()*, *pccard\_lock()*,  
*pccard\_raw\_read()*, *pccard\_unlock()*

## ***pccard\_info()***

© 2004, QNX Software Systems Ltd.

*Obtain socket information from the devp-pccard server*

### **Synopsis:**

```
#include <sys/pccard.h>

int pccard_info(pccard_t handle,
 int socket,
 struct _pccard_info* info,
 unsigned size);
```

### **Arguments:**

|               |                                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>handle</i> | The handle returned by <i>pccard_attach()</i> .                                                                                                                                                               |
| <i>socket</i> | Contains both the socket number as well as the function within the socket. This is achieved by shifting the function number left 8 bits and ORing it with the socket number. The socket number is zero-based. |
| <i>info</i>   | A pointer to a <b>_pccard_info</b> structure that the function fills with the socket information. For more information, see below.                                                                            |
| <i>size</i>   | Size of the <b>_pccard_info</b> structure.                                                                                                                                                                    |

### **Library:**

libpccard

### **Description:**

The *pccard\_info()* function call retrieves socket setup information from the **devp-pccard** server. The information is returned in the **\_pccard\_info** structure.

#### **\_pccard\_info structure**

The **\_pccard\_info** structure is defined in **<pccard.h>** as:

```

struct _pccard_info {
 uint16_t socket; // Socket number (0 based)
 uint16_t status; // Card status (from socket services spec)
 uint32_t flags; // Flags (_PCCARD_FLAG_*)
 uint8_t vcc; // Current Vcc (in tenths of volts)
 uint8_t vpp; // Current Vpp (in tenths of volts)
 uint8_t num_windows; // Number of windows described below
 uint8_t index; // Index for CardBus devices
 uint16_t manufacturer; // Manufacturer ID from PCCARD
 uint16_t card_type; // Card Type from PCCARD
 uint16_t device_id; // CardBus device id
 uint16_t vendor_id; // CardBus vendor id
 uint16_t busnum; // PCI bus number
 uint16_t devfuncnum; // PCI device and function number
 struct _pccard_window {
 uint16_t window; // Window type (_PCCARD_WINDOW_*)
 uint16_t flags; // Window flags (_PCCARD_WINFLAG_*)
 mpid_t pid; // Locking pid
 uint16_t device; // Device type (_PCCARD_DEV_*)
 uint16_t dummy;
 uint32_t dev_size; // Size of memory device
 uint32_t reserved3;
 union {
 struct _pccard_irq {
 uint16_t flags; // (_PCCARD_IRQFLAG_*)
 uint16_t irq;
 }
 struct _pccard_memio {
 uint32_t base; // Base address (in host address space)
 uint32_t size; // Size of window
 uint32_t offset; // offset of region from base of card
 uint16_t flags; // (_PCCARD_MEMIOFLAG_*)
 uint16_t dummy2;
 } memio;
 } un;
 } window[_PCCARD_MAX_WINDOWS];
};

```

**Returns:**

- A positive integer
  - Success. The *socket* parameter is returned.
- 1 An error occurred (*errno* is set).

**Errors:**

ENODEV      Invalid *socket* parameter.

**Examples:**

See *pccard\_arm()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_attach()*, *pccard\_detach()*, *pccard\_lock()*,  
*pccard\_raw\_read()*, *pccard\_unlock()*



## Synopsis:

```
#include <sys/pccard.h>

int pccard_lock(pccard_t handle,
 int socket,
 int index,
 int oflag);
```

## Arguments:

|               |                                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>handle</i> | The handle returned by <i>pccard_attach()</i> .                                                                                                                                                               |
| <i>socket</i> | Contains both the socket number as well as the function within the socket. This is achieved by shifting the function number left 8 bits and ORing it with the socket number. The socket number is zero-based. |
| <i>index</i>  | The window/function number that you want to lock. You can get the window number from the <code>_pccard_info</code> structure (see <i>pccard_info()</i> ).                                                     |
| <i>oflag</i>  | Created by ORing the values required (e.g. <code>O_RDWR   O_EXCL</code> ) for read/write and exclusive access.                                                                                                |

## Library:

`libpccard`

## Description:

The *pccard\_lock()* function call provides exclusive or shared access to the PC Card in *socket* and also sets access permissions.

## Returns:

A positive integer  
Success.

-1 An error occurred (*errno* is set).

**Errors:**

|        |                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------|
| EBADF  | Invalid <i>handle</i> parameter.                                                                         |
| EBUSY  | The window is already locked by another process.                                                         |
| ENODEV | Invalid <i>socket</i> parameter, no PC Card is present in the socket, or invalid <i>index</i> parameter. |

**Examples:**

See *pccard\_arm()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_attach()*, *pccard\_detach()*, *pccard\_info()*,  
*pccard\_raw\_read()*, *pccard\_unlock()*

## Synopsis:

```
#include <sys/pccard.h>

ssize_t pccard_raw_read(pccard_t handle,
 int socket,
 int type,
 unsigned addr,
 ssize_t len,
 void* buf);
```

## Arguments:

|               |                                                                                                                                                                                                               |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>handle</i> | The handle returned by <i>pccard_attach()</i> .                                                                                                                                                               |
| <i>socket</i> | Contains both the socket number as well as the function within the socket. This is achieved by shifting the function number left 8 bits and ORing it with the socket number. The socket number is zero-based. |
| <i>type</i>   | The type of memory that you want to read. Valid values are: <ul style="list-style-type: none"><li>• <code>_PCCARD_MEMTYPE_COMMON</code></li><li>• <code>_PCCARD_MEMTYPE_ATTRIBUTE</code></li></ul>            |
| <i>addr</i>   | The memory address that you want to read from the CIS.                                                                                                                                                        |
| <i>len</i>    | The size of the memory that you want to read.                                                                                                                                                                 |
| <i>buf</i>    | A pointer to a buffer where the function can store the information that it reads from the PC Card.                                                                                                            |

## Library:

`libpccard`

**Description:**

The *pccard\_raw\_read()* function returns the raw CIS (Card Information Structure) data from the PC Card.

**Returns:**

A positive integer

Success. The length read is returned.

-1 An error occurred (*errno* is set).

**Errors:**

EBADF Invalid *handle* parameter.

ENODEV Invalid *socket* parameter.

**Examples:**

See *pccard\_arm()*.

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_attach()*, *pccard\_detach()*, *pccard\_info()*,  
*pccard\_lock()*, *pccard\_unlock()*

## Synopsis:

```
#include <sys/pccard.h>

int pccard_unlock(pccard_t handle,
 int socket,
 int index);
```

## Arguments:

*handle*     The handle returned by *pccard\_attach()*.

*socket*     Contains both the socket number as well as the function within the socket. This is achieved by shifting the function number left 8 bits and ORing it with the socket number. The socket number is zero-based.

*index*     The window/function number that you want to unlock. You can get the window number from the `_pccard_info` structure (see *pccard\_info()*).

## Library:

`libpccard`

## Description:

The *pccard\_unlock()* function unlocks a window previously locked by a call to *pccard\_lock()*. It can only unlock a window locked by the same process ID — you can't unlock a window locked by another process.

## Returns:

0     Success.

-1     An error occurred (*errno* is set).

**Errors:**

- EBADF      Invalid *handle* parameter.
- ENODEV    Invalid *socket* parameter, no PC Card is present in the socket, or invalid *index* parameter.

**Examples:**

See *pccard\_arm()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pccard\_arm()*, *pccard\_attach()*, *pccard\_detach()*, *pccard\_info()*,  
*pccard\_lock()*, *pccard\_raw\_read()*

## Synopsis:

```
#include <hw/pci.h>

int pci_attach(unsigned flags);
```

## Arguments:

*flags*     There are currently no flags defined for this function.

## Library:

`libc`

## Description:

The *pci\_attach()* function connects to the Peripheral Component Interconnect (PCI) server.



You must call *pci\_attach()* before calling any of the other PCI functions.

---

## Returns:

A handle used for calling *pci\_detach()*, or `-1` if an error occurs.

## Errors:

See *open()*.

## Classification:

QNX Neutrino

---

### Safety

Cancellation point    Yes

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config16()*, *pci\_write\_config32()*

“Peripheral Component Interconnect (**pci** - \*)” in the Utilities  
Summary chapter of the *Utilities Reference*



### **Synopsis:**

```
#include <hw/pci.h>

void* pci_attach_device(
 void* handle,
 uint32_t flags,
 uint16_t idx,
 struct pci_dev_info* info);
```

### **Arguments:**

- handle* A handle that identifies the device. The first time you call this function, set *handle* to NULL. This function returns a handle that you can use in a subsequent call to allocate resources for the device.
- flags* Flags that tell the PCI server how you want it to handle resources, which resources to scan for, and which resources to allocate; see “Flags,” below.
- idx* The index of the device: 0 for the first device, 1 for the second, and so on.
- info* A pointer to a `pci_dev_info` structure that specifies the class code, vendor/device ID, or bus number and device/function number that you want to scan for. The function fills in this structure with information about the device.

### **Library:**

`libc`

### **Description:**

The `pci_attach_device()` function attaches a driver to a PCI device.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

Typically drivers use this function to attach themselves to a PCI device, so that other drivers can't attach to the same device. If you specify the `PCLSHARE` flag (see "Flags," below), then multiple drivers can attach to the same device.

The server can scan based on a class code, vendor/device ID, or bus number and device/function number. To control the server scanning, initialize the appropriate fields of the *info* structure and set the appropriate flags.

When you first attach to an uninitialized device, the PCI server assigns all the I/O ports, memory and IRQs required for the device. It also does the IRQ routing for you. Once this has completed successfully, it fills in all these values into your `pci_dev_info` structure to return these values to your application.

When a driver attaches to a device, the PCI server allocates the necessary resources for the device from `procnto` using the *rsrddbmgr\** calls. On X86 BIOS systems, these resources are normally allocated by the BIOS, but on non-x86 systems, these resources have to be allocated from `procnto`.

You can detach the device by passing its handle to *pci\_detach\_device()*. If you call *pci\_detach()*, any resources that *pci\_attach\_device()* allocates are freed.

## **pci\_dev\_info structure**

This function fills in a `pci_dev_info` structure that describes an occurrence of a device.



---

The *pci\_attach\_device()* function doesn't map any of the I/O or memory regions into the process's address space. The addresses returned in the `pci_dev_info` structure are all physical addresses.

---

This structure has the following members:

`uint16_t DeviceId`

The device ID (input/output).

`uint16_t VendorId`

The vendor ID (input/output).

`uint16_t SubsystemId`

The subsystem ID (output).

`uint16_t SubsystemVendorId`

The subsystem vendor ID (output).

`uint8_t BusNumber`

The bus number (output).

`uint8_t DevFunc`

The device/function number (output).

`uint8_t Revision`

The device revision (output).

`uint32_t Class` The class code (input/output).

`uint32_t Irq` The interrupt number (output).

`uint64_t CpuIoTranslation`

The CPU-to-PCI translation value (*pci\_addr = cpu\_addr - translation*).

`uint64_t CpuMemTranslation`

The CPU-to-PCI memory translation (*pci\_addr = cpu\_addr - translation*).

**uint64\_t** *CpuIsaTranslation*

The CPU-to-ISA memory translation ( $pci\_addr = cpu\_addr - translation$ ).

**uint64\_t** *CpuBmstrTranslation*

The translation from the CPU busmaster address to the PCI busmaster address ( $pci\_addr = cpu\_addr + translation$ ).

**uint64\_t** *PciBaseAddress* [6]

The PCI base address (array of six **uint64\_t** items).



---

This function decodes bits 1 and 2 to see whether the register is 32 or 64 bits wide, hence the 64-bit values for the base registers.

---

**uint64\_t** *CpuBaseAddress* [6]

The CPU base address (an array of six **uint64\_t** items).

Some platforms translate addresses across PCI bridges, so that there's one address on the PCI side of the bridge and another on the CPU side. Under x86, the *PciBaseAddress* and *CpuBaseAddress* are the same, but under other platforms, these will be different. In your user application you should always use the *CpuBaseAddress*.

**uint32\_t** *BaseAddressSize* [6]

The size of the base address aperture into the board (an array of six **uint32\_t** items).

**uint64\_t** *PciRom*

The PCI ROM address.

**uint64\_t** *CpuRom*

The CPU ROM address.

`uint32_t RomSize`

The size of the aperture into the board.

## Flags

The *flags* parameter tells the PCI server how resources are to be handled, which resources to scan for, and which resources to allocate.

These bits control how resources are handled:

`PCI_SHARE` Allow resources to be shared with other drivers. If this isn't set, no other driver can attach to the device.

`PCI_PERSIST` Resources persist after the device is detached.

The following bits ask the PCI server to scan for a device based on the fields that you specified in the structure pointed to by *info*:

`PCI_SEARCH_VEND`

*VendorID*

`PCI_SEARCH_VENDEV`

*DeviceId* and *VendorId*

`PCI_SEARCH_CLASS`

*Class*

`PCI_SEARCH_BUSDEV`

*BusNumber* and *DevFunc*

These bits specify which members of the structure the server should initialize:

`PCI_INIT_IRQ` *Irq*

`PCI_INIT_ROM` *PciRom* and *CpuRom*

PCI\_INIT\_BASE0 ... PCI\_INIT\_BASE5

The specified entries of the *PciBaseAddress* and *CpuBaseAddress* arrays

PCI\_INIT\_ALL All members except *PciRom* and *CpuRom*

If you pass 0 for the flags, the default is PCI\_SEARCH\_VENDEV.

### Testing and converting addresses

To facilitate the testing of addresses returned by the PCI server, at least the following macros are defined in the `<pci.h>` header file:

*PCI\_IS\_IO( address )*

Test whether the address is an I/O address.

*PCI\_IS\_MEM( address )*

Test whether the address is a memory address.

*PCI\_IO\_ADDR( address )*

Convert the address returned by the PCI server to an I/O address.

*PCI\_MEM\_ADDR( address )*

Convert the address returned by the PCI server to a memory address.

*PCI\_ROM\_ADDR( address )*

Convert the address returned by the PCI server to a ROM address.

For example:

```
{
 uint64_t port;

 /* Test the address returned by the pci server */
 if (PCI_IS_IO(addr))
 port = (PCI_IO_ADDR(addr));
}
```

**Returns:**

A handle to be used for other *pci\_\** calls associated with a handle, or NULL if an error occurs (*errno* is set).

**Errors:**

- |        |                                                                                                                                   |
|--------|-----------------------------------------------------------------------------------------------------------------------------------|
| EBUSY  | An application has already attached to the device. If it's safe to share the device, specify PCI_SHARE in the <i>flags</i> field. |
| EINVAL | The function couldn't attach a resource to the device.                                                                            |
| ENODEV | This device wasn't found.                                                                                                         |

**Examples:**

Attach to and allocate all resources for the first occurrence of an Adaptec 2940 adapter:

```
#include <hw/pci.h>
#include <hw/pci_devices.h>
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int pidx;
 void* hdl;
 int phdl;
 struct pci_dev_info inf;

 /* Connect to the PCI server */
 phdl = pci_attach(0);
 if(phdl == -1) {
 fprintf(stderr, "Unable to initialize PCI\n");

 return EXIT_FAILURE;
 }

 /* Initialize the pci_dev_info structure */
 memset(&inf, 0, sizeof(inf));
 pidx = 0;
 inf.VendorId = PCI_VENDOR_ID_ADAPTEC;
 inf.DeviceId = PCI_DEVICE_ID_ADAPTEC_2940F;
}
```

```
hdl = pci_attach_device(NULL, PCI_INIT_ALL, pidx, &inf);
if(hdl == NULL) {
 fprintf(stderr, "Unable to locate adapter\n");
} else {
 /* Do something to the adapter */
 pci_detach_device(hdl);
}

/* Disconnect from the PCI server */
pci_detach(phdl);

return EXIT_SUCCESS;
}
```

Attach to the first occurrence of an Adapter 2940 adapter and allocate resources in a second call:

```
#include <hw/pci.h>
#include <hw/pci_devices.h>
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int pidx;
 void* hdl;
 void* retval;
 int phdl;
 struct pci_dev_info inf;

 phdl = pci_attach(0);
 if(phdl == -1) {
 fprintf(stderr, "Unable to initialize PCI\n");

 return EXIT_FAILURE;
 }

 memset(&inf, 0, sizeof(inf));
 pidx = 0;
 inf.VendorId = PCI_VENDOR_ID_ADAPTEC;
 inf.DeviceId = PCI_DEVICE_ID_ADAPTEC_2940F;

 hdl = pci_attach_device(NULL, 0, pidx, &inf);
 if(hdl == NULL) {
 fprintf(stderr, "Unable to locate adapter\n");
 }

 retval = pci_attach_device(hdl, PCI_INIT_ALL, pidx, &inf);
 if(retval == NULL) {
```



```
 fprintf(stderr, "Unable allocate resources\n");
 }

 pci_detach(phdl);

 return EXIT_SUCCESS;
}
```

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pci\_attach()*, *pci\_detach()*, *pci\_detach\_device()*, *pci\_find\_class()*,  
*pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*, *pci\_read\_config8()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*

## ***pci\_detach()***

© 2004, QNX Software Systems Ltd.

*Disconnect from the PCI server*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_detach(unsigned handle);
```

### **Arguments:**

*handle*     The value returned by a successful call to *pci\_attach()*.

### **Library:**

`libc`

### **Description:**

The *pci\_detach()* function disconnects from the PCI server. Any resources allocated with *pci\_attach\_device()* are released.

The *pci\_attach()* function opens a file descriptor against the PCI server, and all of the low-level library calls to the PCI server use this fd. When you call *pci\_detach()*, the low-level code does a *close()* on the file descriptor, which tells the PCI server to clean up any allocations associated with it.



---

Don't call any of the other *pci\_\**(*)* functions after calling *pci\_detach()* (unless you've reattached with *pci\_attach()*).

---

### **Returns:**

`PCL_SUCCESS`.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach(), pci\_attach\_device(), pci\_detach\_device(), pci\_find\_class(), pci\_find\_device(), pci\_present(), pci\_read\_config(), pci\_read\_config8(), pci\_read\_config16(), pci\_read\_config32(), pci\_rescan\_bus(), pci\_write\_config(), pci\_write\_config8(), pci\_write\_config16(), pci\_write\_config32()*

## ***pci\_detach\_device()***

© 2004, QNX Software Systems Ltd.

*Detach a driver from a PCI device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_detach_device(void* handle);
```

### **Arguments:**

*handle*     The handle returned by *pci\_attach\_device()*.

### **Library:**

libc

### **Description:**

The *pci\_detach\_device()* function detaches a driver from a PCI device. Any resources allocated with *pci\_attach\_device()* are released, unless you attached the device with the PCLPERSIST flag set.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

### **Returns:**

PCLDEVICE\_NOT\_FOUND

No device could be found for *handle*.

PCLSUCCESS

Success.

-1     You haven't called *pci\_attach()*, or the call to it failed.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_find\_class()*,  
*pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*, *pci\_read\_config8()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_find\_class()***

© 2004, QNX Software Systems Ltd.

*Find devices that have a specific class code*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_find_class(unsigned long class_code,
 unsigned index,
 unsigned* bus,
 unsigned* dev_func);
```

### **Arguments:**

|                   |                                                                                                                                                                                |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>class_code</i> | The class of device or function that you want to find.                                                                                                                         |
| <i>index</i>      | The index of the device or function that you want to find: 0 for the first, 1 for the second, and so on.                                                                       |
| <i>bus</i>        | The bus number, in the range [0...255].                                                                                                                                        |
| <i>dev_func</i>   | The device or function number of the <i>n</i> th device or function of the given class. The device number is in bits 7 through 3, and the function number in bits 2 through 0. |

### **Library:**

libc

### **Description:**

The *pci\_find\_class()* function determines the location of the *n*th PCI device or function that has the specified class code.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

You can find all the devices having the same class code by making successive calls to this function, starting with an index of 0, and incrementing it until PCI\_DEVICE\_NOT\_FOUND is returned.

## Returns:

PCI\_DEVICE\_NOT\_FOUND

The device or function wasn't found.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

## Classification:

QNX Neutrino

### Safety

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes

## See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*, *pci\_read\_config8()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_find\_device()***

© 2004, QNX Software Systems Ltd.

*Find the PCI device with a given device ID and vendor ID*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_find_device(unsigned device,
 unsigned vendor,
 unsigned index,
 unsigned* bus,
 unsigned* dev_func);
```

### **Arguments:**

|                 |                                                                                                                                                                                                                                                             |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>device</i>   | The Device ID.                                                                                                                                                                                                                                              |
| <i>vendor</i>   | The Vendor ID.                                                                                                                                                                                                                                              |
| <i>index</i>    | The index ( <i>n</i> ) of the device or function sought.                                                                                                                                                                                                    |
| <i>bus</i>      | A pointer to a location where the function can store the bus number of the device or function found.                                                                                                                                                        |
| <i>dev_func</i> | A pointer to a location where the function can store the device or function ID of the <i>n</i> th device or function found with the specified device and vendor IDs. The device number is in bits 7 through 3, and the function number in bits 2 through 0. |

### **Library:**

`libc`

### **Description:**

The *pci\_find\_device()* function returns the location of the *n*th PCI device that has the specified Device ID and Vendor ID.





---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

You can find all the devices having the same Device and Vendor IDs by making successive calls to this function, starting with an index of 0, and incrementing it until `PCI_DEVICE_NOT_FOUND` is returned.

### Returns:

`PCI_DEVICE_NOT_FOUND`

The device or function wasn't found.

`PCI_SUCCESS`

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

### Classification:

QNX Neutrino

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_present()*, *pci\_read\_config()*, *pci\_read\_config8()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_irq\_routing\_options()***

© 2004, QNX Software Systems Ltd.

*Retrieve PCI IRQ routing information*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_irq_routing_options (
 IRQRoutingOptionsBuffer * buf,
 uint32_t * irq);
```

### **Arguments:**

- buf* A pointer to a **IRQRoutingOptionsBuffer** structure where the function can store the IRQ routing information. For information about the layout of this buffer, see *PCI BIOS SPECIFICATION* Revision 2.1. You can get it from the PCI Special Interest Group at <http://pcisig.com/>.
- irq* A pointer to a location where the function can store the current state of interrupts.

### **Library:**

**libc**

### **Description:**

The *pci\_irq\_routing\_options()* function returns the following:

- PCI interrupt routing options available on the system motherboard
- the current state of interrupts that are currently exclusively assigned to PCI.

Routing information is returned in a data buffer that contains an IRQ routing for each PCI device or slot.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

PCLSUCCESS      Success.

-1                You haven't called *pci\_attach()*, or the call to it failed.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <errno.h>
#include <inttypes.h>
#include <hw/pci.h>
#include <sys/neutrino.h>

struct {
 IRQRoutingOptionsBuffer buf;
 uint8_t databuf [2048];
} route_buf;

int main (void)
{
 int phdl;
 uint32_t irq;

 if ((phdl = pci_attach (0)) == -1) {
 printf ("Unable to attach - errno %s\n", strerror (errno));
 exit (1);
 }

 memset (route_buf.databuf, 0, sizeof (route_buf.databuf));
 route_buf.buf.BufferSize = sizeof (route_buf.databuf);
 if (pci_irq_routing_options (&route_buf.buf, &irq) != PCISUCCESS) {
 printf ("Routing option failed - errno %s\n", strerror (errno));
 exit (1);
 }

 printf ("PCI Irq Map = %x\n", irq);
 pci_detach (phdl);
 return (0);
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config16()*, *pci\_write\_config32()*

## Synopsis:

```
#include <hw/pci.h>

int pci_map_irq(unsigned bus,
 unsigned dev_func,
 short intno,
 short intpin);
```

## Arguments:

|                 |                                                                                                                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number of the device.                                                                                                          |
| <i>dev_func</i> | The device or function number of the device. The device number is in bits 7 through 3, and the function number is in bits 2 through 0. |
| <i>intno</i>    | The interrupt to be mapped (0 - 15).                                                                                                   |
| <i>intpin</i>   | The PCI interrupt pin (0x0a - 0x0d).                                                                                                   |

## Library:

libc

## Description:

The *pci\_map\_irq()* function maps a PCI interrupt pin to a specific interrupt request (IRQ).



You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

## Returns:

PCL\_SUCCESS

Success.

PCL\_SET\_FAILED

The PCI server was unable to map the *intno/intpin*.

PCI\_UNSUPPORTED\_FUNCTION

This function isn't supported by the BIOS.

-1 You haven't called *pci\_attach()*, or the call to it failed.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config16()*, *pci\_write\_config32()*

## Synopsis:

```
#include <hw/pci.h>

int pci_present(unsigned* lastbus,
 unsigned* version,
 unsigned* hardware);
```

## Arguments:

*lastbus*      The number of the last PCI bus in the system. PCI buses are numbered from 0, up to and including this value.

*version*      The version number of the PCI interface.

*hardware*    The specific hardware characteristics the platform supports with regard to accessing configuration space and generating PCI Special Cycles.

The PCI specification defines two hardware mechanisms for accessing configuration space. Bit 0 of *hardware* is set (1) if mechanism 1 is supported, and reset (0) otherwise. Bit 1 is set (1) if mechanism 2 is supported, and reset (0) otherwise.

The specification also defines hardware mechanisms for generating Special Cycles. Bit 4 of *hardware* is set (1) if the platform supports Special Cycle generation based on Config Mechanism 1, and reset (0) otherwise. Bit 5 is set (1) if the platform supports Special Cycle generation based on Config Mechanism 2, and reset (0) otherwise.

The arguments can be NULL if you just want to check for PCI capabilities.

**Library:**

libc

**Description:**

The *pci\_present()* function determines whether or not the PCI BIOS interface function set is present. It also determines the following:

- the current interface version
- what hardware mechanism for accessing configuration space is supported
- whether or not the hardware supports the generation of PCI Special Cycles.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

-1 PCI BIOS isn't present.

PCLSUCCESS

PCI BIOS is present.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes



**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config16()*, *pci\_write\_config32()*

## ***pci\_read\_config()***

© 2004, QNX Software Systems Ltd.

*Read from the configuration space of a PCI device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_read_config(void* handle,
 unsigned offset,
 unsigned count,
 size_t size,
 void* buff);
```

### **Arguments:**

|               |                                                                               |
|---------------|-------------------------------------------------------------------------------|
| <i>handle</i> | The handle returned by <i>pci_attach_device()</i> .                           |
| <i>offset</i> | The offset into the configuration space where you want to read from.          |
| <i>count</i>  | The number of objects that you want to read.                                  |
| <i>size</i>   | The size of each object.                                                      |
| <i>buff</i>   | A pointer to a buffer where the function can store the objects that it reads. |

### **Library:**

`libc`

### **Description:**

The *pci\_read\_config()* function reads *count* objects of the specified *size* into *buff* at the given *offset* from the configuration space of the PCI device specified by *handle*.



You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

## Returns:

PCI\_BAD\_REGISTER\_NUMBER

The offset is invalid.

PCI\_BUFFER\_TOO\_SMALL

The PCI BIOS server reads only 100 bytes at a time; *size* is too large.

PCI\_DEVICE\_NOT\_FOUND

The *handle* is invalid.

PCISUCCESS

Success.

-1 You haven't called *pci\_attach()*, or the call to it failed.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config8()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_read\_config8()***

© 2004, QNX Software Systems Ltd.

*Read a byte from the configuration space of a device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_read_config8(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### **Arguments:**

|                 |                                                                           |
|-----------------|---------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                           |
| <i>dev_func</i> | The name of the device or function.                                       |
| <i>offset</i>   | The register offset into the configuration space, in the range [0...255]. |
| <i>count</i>    | The number of bytes to read.                                              |
| <i>buff</i>     | A pointer to a buffer where the requested bytes are placed.               |

### **Library:**

libc

### **Description:**

The *pci\_read\_config8()* function reads the specified number of bytes from the configuration space of the given device or function.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

## Returns:

PCI\_BAD\_REGISTER\_NUMBER

An invalid register offset was given.

PCI\_BUFFER\_TOO\_SMALL

The PCI BIOS server reads only 100 bytes at a time; *count* is too large.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

## Classification:

QNX Neutrino

### Safety

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes

## See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config16()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_read\_config16()***

© 2004, QNX Software Systems Ltd.

*Read 16-bit values from the configuration space of a device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_read_config16(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### **Arguments:**

|                 |                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                                                                                       |
| <i>dev_func</i> | The name of the device or function.                                                                                                   |
| <i>offset</i>   | The register offset into the configuration space. This offset must be aligned to a 16-bit boundary (that is 0, 2, 4, ..., 254 bytes). |
| <i>count</i>    | The number of 16-bit values to read.                                                                                                  |
| <i>buff</i>     | A pointer to a buffer where the requested 16-bit values are placed.                                                                   |

### **Library:**

`libc`

### **Description:**

The *pci\_read\_config16()* function reads the specified number of 16-bit values from the configuration space of the given device or function.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

## Returns:

PCI\_BAD\_REGISTER\_NUMBER

An invalid offset register number was given.

PCI\_BUFFER\_TOO\_SMALL

The PCI BIOS server reads only 50 words at a time; *count* is too large.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config32()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_read\_config32()***

© 2004, QNX Software Systems Ltd.

*Read 32-bit values from the configuration space of a device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_read_config32(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### **Arguments:**

|                 |                                                                                                                                       |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                                                                                       |
| <i>dev_func</i> | The name of the device or function.                                                                                                   |
| <i>offset</i>   | The register offset into the configuration space. This offset must be aligned to a 32-bit boundary (that is 0, 4, 8, ..., 252 bytes). |
| <i>count</i>    | The number of 32-bit values to read.                                                                                                  |
| <i>buff</i>     | A pointer to a buffer where the requested 32-bit values are placed.                                                                   |

### **Library:**

`libc`

### **Description:**

The *pci\_read\_config32()* function reads the specified number of 32-bit values from the configuration space of the given device or function.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---



## Returns:

PCI\_BAD\_REGISTER\_NUMBER

An invalid register offset was given.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_rescan\_bus()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_rescan\_bus()***

© 2004, QNX Software Systems Ltd.

*Rescan the PCI bus for added or removed devices*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_rescan_bus(void);
```

### **Library:**

libc

### **Description:**

The *pci\_rescan\_bus()* function asks the PCI server to rescan the PCI bus(es) for devices that have been inserted or removed. This is used in hot swap situations such as for CardBus cards. The PCI server updates its internal configuration to reflect any changes.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

### **Returns:**

PCLSUCCESS  
Success.

-1 The function failed.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_write\_config()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

## ***pci\_write\_config()***

© 2004, QNX Software Systems Ltd.

*Write to the configuration space of a PCI device*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_write_config(void* handle,
 unsigned offset,
 unsigned count,
 size_t size,
 const void* buff);
```

### **Arguments:**

|               |                                                                           |
|---------------|---------------------------------------------------------------------------|
| <i>handle</i> | The handle returned by <i>pci_attach_device()</i> .                       |
| <i>offset</i> | The offset into the configuration space where you want to write the data. |
| <i>count</i>  | The number of objects that you want to write.                             |
| <i>size</i>   | The size of each object.                                                  |
| <i>buff</i>   | A pointer to the data that you want to write.                             |

### **Library:**

`libc`

### **Description:**

The *pci\_write\_config()* function writes *count* objects of the specified *size* from *buff* at the given *offset* to the configuration space of the PCI device specified by *handle*.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

PCI\_BAD\_REGISTER\_NUMBER

The *offset* specified is invalid.

PCI\_BUFFER\_TOO\_SMALL

The *size* argument is too large.

PCI\_SET\_FAILED

An error occurred writing to the configuration space of the device.

PCI\_SUCCESS

Success.

PCI\_UNSUPPORTED\_FUNCT

This device doesn't support writing to its configuration space.

-1 You haven't called *pci\_attach()*, or the call to it failed.**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,

## *pci\_write\_config()*

---

© 2004, QNX Software Systems Ltd.

*pci\_rescan\_bus()*, *pci\_write\_config8()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*

### **Synopsis:**

```
#include <hw/pci.h>

int pci_write_config8(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### **Arguments:**

|                 |                                                                                                                   |
|-----------------|-------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                                                                   |
| <i>dev_func</i> | The device or function ID. The device number is in bits 7 through 3, and the function number in bits 2 through 0. |
| <i>offset</i>   | The register offset into the configuration space, in the range [0...255].                                         |
| <i>count</i>    | The number of bytes to write.                                                                                     |
| <i>buff</i>     | A pointer to a buffer containing the data to be written into the configuration space.                             |

### **Library:**

`libc`

### **Description:**

The *pci\_write\_config8()* function writes individual bytes to the configuration space of the specified device.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

PCI\_BAD\_REGISTER\_NUMBER

An invalid offset register number was given.

PCI\_BUFFER\_TOO\_SMALL

The *size* argument is greater than 100 bytes.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config16()*,  
*pci\_write\_config32()*



### Synopsis:

```
#include <hw/pci.h>

int pci_write_config16(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### Arguments:

|                 |                                                                                                                       |
|-----------------|-----------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                                                                       |
| <i>dev_func</i> | The device or function ID. The device number is in bits 7 through 3, and the function number in bits 2 through 0.     |
| <i>offset</i>   | The offset into the configuration space. This must be aligned to a 16-bit boundary (that is 0, 2, 4, ..., 254 bytes). |
| <i>count</i>    | The number of 16-bit values to write.                                                                                 |
| <i>buff</i>     | A pointer to a buffer containing the data to be written into the configuration space.                                 |

### Library:

`libc`

### Description:

The *pci\_write\_config16()* function writes individual 16-bit values to the configuration space of the specified device.



You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

PCI\_BAD\_REGISTER\_NUMBER

An invalid register offset was given.

PCI\_BUFFER\_TOO\_SMALL

The *size* argument is greater than 50 words.

PCI\_SUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config32()*

### Synopsis:

```
#include <hw/pci.h>

int pci_write_config32(unsigned bus,
 unsigned dev_func,
 unsigned offset,
 unsigned count,
 char* buff);
```

### Arguments:

|                 |                                                                                                                                |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------|
| <i>bus</i>      | The bus number.                                                                                                                |
| <i>dev_func</i> | The device or function ID. The device number is in bits 7 through 3, and the function number in bits 2 through 0.              |
| <i>offset</i>   | The register offset into the configuration space. This must be aligned to a 32-bit boundary (that is 0, 4, 8, ..., 252 bytes). |
| <i>count</i>    | The number of 32-bit values to write.                                                                                          |
| <i>buff</i>     | A pointer to a buffer containing the data to be written into the configuration space.                                          |

### Library:

`libc`

### Description:

The *pci\_write\_config32()* function writes individual 32-bit values to the configuration space of the specified device.



---

You must successfully call *pci\_attach()* before calling any of the other PCI functions.

---

**Returns:**

PCI\_BAD\_REGISTER\_NUMBER

An invalid register offset was given.

PCLSUCCESS

The device or function was found.

-1 You haven't called *pci\_attach()*, or the call to it failed.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point Yes

Interrupt handler No

Signal handler Yes

Thread Yes

**See also:**

*pci\_attach()*, *pci\_attach\_device()*, *pci\_detach()*, *pci\_detach\_device()*,  
*pci\_find\_class()*, *pci\_find\_device()*, *pci\_present()*, *pci\_read\_config()*,  
*pci\_read\_config8()*, *pci\_read\_config16()*, *pci\_read\_config32()*,  
*pci\_rescan\_bus()*, *pci\_write\_config()*, *pci\_write\_config8()*,  
*pci\_write\_config16()*

## Synopsis:

```
#include <stdio.h>

int pclose(FILE* stream);
```

## Arguments:

*stream* The stream pointer for the pipe that you want to close, that you obtained by calling *popen()*.

## Library:

libc

## Description:

The *pclose()* function closes the pipe associated with *stream*, and waits for the subprocess created by *popen()* to terminate.

## Returns:

The termination status of the command language interpreter, or -1 if an error occurred (*errno* is set).

## Errors:

EINTR The *pclose()* function was interrupted by a signal while waiting for the child process to terminate.

ECHILD The *pclose()* function was unable to obtain the termination status of the child process.

## Examples:

See *popen()*.

**Classification:**

POSIX 1003.1a

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*errno, popen(), pipe()*

## Synopsis:

```
#include <stdio.h>

void perror(const char *prefix);
```

## Arguments:

*prefix*     NULL, or a string that you want to print before the error message.

## Library:

libc

## Description:

The *perror()* function prints the following to *stderr*:

- the given *prefix*, followed by “: ”
- the error message returned by *strerror()* for the current value of *errno*
- a newline character.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 FILE *fp;

 fp = fopen("data.fil", "r");
 if(fp == NULL) {
 perror("Unable to open file");
 return EXIT_FAILURE;
 }
 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, fprintf(), stderr, strerror()*



## Synopsis:

```
#include <unistd.h>

int pipe(int fildes[2]);
```

## Arguments:

*fildes* An array where the function can store the file descriptors for the ends of the pipe.

## Library:

libc

## Description:

The *pipe()* function creates a pipe (an unnamed FIFO) and places a file descriptor for the read end of the pipe in *fildes*[0], and a file descriptor for the write end of the pipe in *fildes*[1]. Their integer values are the two lowest available at the time of the *pipe()* function call. The O\_NONBLOCK flag is cleared for both file descriptors. (You can use *fcntl()* to set the O\_NONBLOCK flag.)

You can write data to file descriptor *fildes*[1] and read it from file descriptor *fildes*[0]. If you read from file descriptor *fildes*[0], it returns the data written to *fildes*[1] on a first-in-first-out (FIFO) basis.

The pipe buffer is allocated by the **pipe** resource manager.

You typically use this function to connect standard utilities acting as filters, passing the write end of the pipe to the data-producing process as its STDOUT\_FILENO, and the read end of the pipe to the data-consuming process as its STDIN\_FILENO (either via the traditional *fork()*, *dup2()*, or *exec\**, or the more efficient *spawn\** calls).

If successful, *pipe()* marks the *st\_mtime*, *st\_ctime*, *st\_atime* and *st\_mtime* fields of the pipe for updating.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EMFILE The calling process doesn't have at least 2 unused file descriptors available.
- ENFILE The number of simultaneously open files in the system would exceed the configured limit.
- ENOSPC There's insufficient space available to allocate the pipe buffer.
- ENOSYS There's no pipe manager running.
- EROFS The pipe pathname space is a read-only filesystem.

**Classification:**

POSIX 1003.1

**Safety**

---

- Cancellation point No
- Interrupt handler No
- Signal handler Yes
- Thread Yes

**See also:**

*errno*, *fcntl()*, *nbaconnect()*, *open()*, *popen()*, *read()*, *write()*  
**pipe** in the *Utilities Reference*

**Synopsis:**

```
#include <sys/poll.h>
int poll(struct pollfd fds*, nfds_t nfds, int timeout)
```

**Arguments:**

*fds*            The array of interest.

*nfds*           Number of elements in the *fds* array.

*timeout*        Timeout in milliseconds.

**Library:**

```
libc
```

**Description:**

The *poll()* function provides applications with a mechanism for multiplexing input/output over a set of file descriptors.

The *pollfd* structure has the following components:

```
struct pollfd {
 int fd;
 short events;
 short revents;
};
```

For each member of the array pointed to by *fds*, *poll()* examines the given file descriptor for the event(s) specified in *events*. The number of *pollfd* structures in the *fds* array is specified by *nfds*. The array's members are *pollfd* structures within which *fd* specifies an open file descriptor, *events* and *revents* are bitmasks constructed by OR'ing a combination of the following event flags:

**POLLERR**      An error has occurred on the device. This flag is valid only in the *revents* bitmask; it's ignored in the *events* member.

- POLLHUP      The device has been disconnected. This event and POLLOUT are mutually exclusive; a device can never be writable if a hangup has occurred. However, this event and POLLIN, POLLRDNORM, POLLRDBAND, or POLLPRI are not mutually exclusive. If the remote end of a socket is closed, *poll()* indicates a POLLIN event rather than POLLHUP. This flag is valid only in the *revents* bitmask; it's ignored in the *events* member.
- POLLIN        Data other than high-priority data may be read without blocking. This is equivalent to POLLRDNORM | POLLRDBAND.
- POLLNVAL      The specified fd value is invalid. This flag is only valid in the *revents* member; it shall ignored in the *events* member.
- POLLOUT      Normal data may be written without blocking.
- POLLPRI      High-priority data may be read without blocking.
- POLLRDBAND      Priority data may be read without blocking.
- POLLRDNORM    Normal data may be read without blocking.
- POLLWRBAND    Priority data may be written.
- POLLWRNORM    Equivalent to POLLOUT.

The significance and semantics of normal, priority, and high-priority data are file- and device-specific.

If the value of *fd* is less than 0, events are ignored; and *revents* are set to 0 in that entry on return from *poll()*.

In each *pollfd* structure, *poll()* clears the *revents* member, except that where the application requested a report on a condition by setting one of the bits of events listed above, *poll()* sets the corresponding bit in *revents* if the requested condition is true. In addition, *poll()* sets the POLLHUP, POLLERR, and POLLNVAL flag in *revents* if the condition is true, even if the application didn't set the corresponding bit in events.

If none of the defined events occurs on any selected file descriptor, *poll()* waits at least *timeout* milliseconds for an event to occur on any of the selected file descriptors. If the value of *timeout* is 0, *poll()* returns immediately. If the value of *timeout* is -1, *poll()* blocks until a requested event occurs or until the call is interrupted.

The *poll()* function isn't affected by the O\_NONBLOCK flag.

The *poll()* function reports regular files, terminal and pseudo-terminal devices, FIFOs, and pipes.

Regular files always poll TRUE for reading and writing.

A file descriptor for a socket that's listening for connections indicates that it's ready for reading, once connections are available. A file descriptor for a socket that connects asynchronously indicates that it's ready for writing, once a connection has been established.

## Returns:

- > 0 Total number of file descriptors that have been selected.
- 0 The call timed out, and no file descriptor has been selected.
- 1 Failure, and *errno* is set.

## Errors:

- EAGAIN The allocation of internal data structures failed but a subsequent request may succeed.
- EINTR A signal was caught during *poll()*
- EFAULT The *fds* argument pointed to a nonexistent portion of the calling process's address space.

**Examples:**

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <sys/poll.h>
#include <arpa/inet.h>
#include <unistd.h>
#include <stdio.h>
#include <pthread.h>
#include <fcntl.h>
#include <errno.h>
#include <string.h>

struct sockaddr_in sad;

void *
client(void *arg)
{
 int s;
 const char *p = "Some data\n";

 if ((s = socket(AF_INET, SOCK_STREAM, 0)) == -1) {
 perror("socket");
 return NULL;
 }

 if (connect(s, (struct sockaddr *)&sad, sizeof(sad)) == -1) {
 perror("connect");
 return NULL;
 }

 write(s, p, strlen(p));
 close(s);

 return NULL;
}

int
main(void)
{
 struct pollfd fds;
 int s = -1, s2 = -1, done_accept = 0, oflags, ret;
 char buf[100];

 if ((s = socket(AF_INET, SOCK_STREAM, 0)) == -1) {
 perror("socket");
 return 1;
 }

 sad.sin_family = AF_INET;
 sad.sin_len = sizeof(sad);
```

```
sad.sin_addr.s_addr = inet_addr("127.0.0.1");
sad.sin_port = htons(1234);

fds.fd = s;
fds.events = POLLRDNORM;

oflags = fcntl(s, F_GETFL);
oflags |= O_NONBLOCK;
fcntl(s, F_SETFL, oflags);

if (bind(s, (struct sockaddr *)&sad, sizeof(sad)) == -1) {
 perror("bind");
 return 1;
}

listen(s, 5);

if ((ret = pthread_create(NULL, NULL, client, NULL)) != EOK) {
 fprintf(stderr, "pthread_create: %s\n", strerror(ret));
 return 1;
}

for (;;) {
 if ((ret = poll(&fds, 1, -1)) == -1) {
 perror("poll");
 break;
 }

 else if (ret != 1 || (fds.revents & POLLRDNORM) == 0) {
 break;
 }

 if (done_accept) {
 if ((ret = read(s2, buf, sizeof(buf))) <= 0) {
 break;
 }

 printf("%s", buf);
 }

 else {
 if ((s2 = accept(s, NULL, 0)) == -1) {
 perror("accept");
 break;
 }

 fds.fd = s2;
 done_accept = 1;
 }
}
```

```
 close(s);
 close(s2);

 return 0;
}
```

### Classification:

POSIX 1003.1

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### Caveats:

Not all managers support POLLPRI, POLLPRI, POLLERR, and POLLHUP.

### See also:

*read()*, *select()*, *write()*

<sys/poll.h>



## Synopsis:

```
#include <stdio.h>

FILE* popen(const char* command,
 const char* mode);
```

## Arguments:

*command*    The command that you want to execute.

*mode*        The I/O mode for the pipe, which must be "r" or "w"; see below.

## Library:

libc

## Description:

The *popen()* function executes the command specified by *command* and creates a pipe between the calling process and the executed command.

Depending on the *mode* argument, you can use the returned stream pointer to read from or write to the pipe.

The executed command has the same environment as its parents. The command is started as follows:

```
spawnlp (P_NOWAIT, shell_command, shell_command,
 "-c", command, (char*)NULL);
```

where *shell\_command* is the command specified by the **SHELL** environment variable (if it exists), or the **sh** utility.

The *mode* argument to *popen()* is a string that specifies an I/O mode for the pipe:

- If *mode* is "r", then when the child process is started:

- Its file descriptor, `STDOUT_FILENO`, is the writable end of the pipe.
- The `fileno( stream )` in the calling process is the readable end of the pipe, where `stream` is the stream pointer returned by `popen()`.
- If `mode` is "`w`", then when the child process is started:
  - Its file descriptor, `STDIN_FILENO`, is the readable end of the pipe.
  - The `fileno( stream )` in the calling process is the writable end of the pipe, where `stream` is the stream pointer return by `popen()`.
- If `mode` is any other value, the result is undefined.



---

Use `pclose()` to close a stream that you used `popen()` to open.

---

## Returns:

A non-NULL stream pointer on successful completion. If `popen()` is unable to create either the pipe or the subprocess, it returns a NULL stream pointer and sets `errno`.

## Errors:

`EINVAL`      The `mode` argument is invalid.  
`ENOSYS`      There's no pipe manager running.

The `popen()` function may also set `errno` values as described by the `pipe()` and `spawnl()` functions.

## Examples:

```
/*
 * upper: executes a given program, converting all input
 * to upper case.
 */
#include <stdio.h>
#include <stdlib.h>
```

```
#include <string.h>
#include <ctype.h>
#include <unistd.h>
#include <limits.h>

char buffer[_POSIX_ARG_MAX];

int main(int argc, char** argv)
{
 int i;
 int c;
 FILE* f;

 for(i = 1; i < argc; i++) {
 strcat(buffer, argv[i]);
 strcat(buffer, " ");
 }
 if((f = popen(buffer, "w")) == NULL) {
 perror("popen");
 return EXIT_FAILURE;
 }
 while((c = getchar()) != EOF) {
 if(islower(c))
 c = toupper(c);
 putc(c, f);
 }
 pclose(f);
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1a

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*errno, pclose(), pipe(), spawnlp()*

## ***posix\_mem\_offset()*, *posix\_mem\_offset64()***

*Get the offset of a mapped typed memory block*

### **Synopsis:**

```
#include <sys/mman.h>

int posix_mem_offset(const void * addr,
 size_t length,
 off_t * offset,
 size_t * contig_len,
 int * fd);

int posix_mem_offset64(const void * addr,
 size_t length,
 off64_t * offset,
 size_t * contig_len,
 int * fd);
```

### **Library:**

libc

### **Description:**



---

The *posix\_mem\_offset()* and *posix\_mem\_offset64()* functions aren't currently supported.

---

The *posix\_mem\_offset()* and *posix\_mem\_offset64()* functions set the variable pointed to by *offset* to the offset (or location), within a typed memory object, of the memory block currently mapped at *addr*.

### **Returns:**

-1 ( *errno* is set).

### **Errors:**

ENOSYS      The *posix\_mem\_offset()* function isn't supported by this implementation.

**Classification:**

*posix\_mem\_offset()* is POSIX 1003.1j (draft); *posix\_mem\_offset64()* is for large-file support

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*mem\_offset()*, *mem\_offset64()*, *mmap()*

## Synopsis:

```
#include <stdlib.h>

int posix_memalign(void ** memptr,
 size_t alignment,
 size_t size);
```

## Arguments:

|                  |                                                                                               |
|------------------|-----------------------------------------------------------------------------------------------|
| <i>memptr</i>    | A pointer to a location where <i>posix_memalign()</i> can store a pointer the memory.         |
| <i>alignment</i> | The alignment to use for the memory. This must be a multiple of <code>size( void * )</code> . |
| <i>size</i>      | The size, in bytes, of the block to allocate.                                                 |

## Library:

`libc`

## Description:

The *posix\_memalign()* function allocates *size* bytes aligned on a boundary specified by *alignment*. It returns a pointer to the allocated memory in *memptr*.

The buffer allocated by *posix\_memalign()* is contiguous in virtual address space, but not physical memory. Since some platforms don't allocate memory in 4K page sizes, you shouldn't assume that the memory allocated will be physically contiguous if you specify a size of 4K or less.

You can obtain the physical address of the start of the buffer using *mem\_offset()* with *fd=NOFD*.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EINVAL The value of *alignment* isn't a multiple of **size (void \* )**.
- ENOMEM There's insufficient memory available with the requested alignment.

**Classification:**

POSIX 1003.1-2001

**Safety**

---

- |                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*, *free()*, *malloc()*, *memalign()*



## Synopsis:

```
#include <math.h>

double pow(double x,
 double y);

float powf(float x,
 float y);
```

## Arguments:

- x* The number you want to raise.
- y* The power you want to raise the number to.

## Library:

`libm`

## Description:

The *pow()* and *powf()* functions compute *x* raised to the power of *y*.

A domain error occurs if *x* = 0, and *y* ≤ 0, or if *x* is negative, and *y* isn't an integer. A range error may also occur.

## Returns:

The value of  $x^y$ .



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
{
 printf("%f\n", pow(1.5, 2.5));

 return EXIT_SUCCESS;
}
```

produces the output:

2.755676

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno*, *exp()*, *log()*, *sqrt()*

### **Synopsis:**

```
#include <unistd.h>

ssize_t pread(int filedes,
 void *buff,
 size_t nbytes,
 off_t offset);

ssize_t pread64(int filedes,
 void *buff,
 size_t nbytes,
 off64_t offset);
```

### **Arguments:**

- filedes*    The descriptor of the file that you want to read from.
- buff*        A pointer to a buffer where the function can store the data that it reads.
- nbytes*      The number of bytes that you want to read.
- offset*      The desired position inside the file.

### **Library:**

`libc`

### **Description:**

The *pread()* function performs the same action as *read()*, except that it reads from a given position in the file without changing the file pointer.

The *pread()* function reads up to the maximum offset value that can be represented in an `off_t` for regular files. An attempt to perform a *pread()* on a file that's incapable of seeking results in an error.

The *pread64()* function is a 64-bit version of *pread()*.

**Returns:**

The number of bytes actually read, or -1 if an error occurred (*errno* is set).

**Errors:**

- EAGAIN     The O\_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the read operation.
- EBADF     The file descriptor, *fdes*, isn't a valid file descriptor open for reading.
- EINTR     The read operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file does not report partial transfers.
- EIO        A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.
- ENOSYS    The *pread()* function isn't implemented for the filesystem specified by *filedes*.

**Classification:**

*pread()* is standard Unix; *pread64()* is for large-file support

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close()*, *creat()*, *dup()*, *dup2()*, *errno*, *fcntl()*, *lseek()*, *open()*, *pipe()*,  
*pwrite()*, *read()*, *readblock()*, *readv()*, *select()*, *write()*, *writeblock()*,  
*writenv()*

# ***printf()***

© 2004, QNX Software Systems Ltd.

*Write formatted output to stdout*

## **Synopsis:**

```
#include <stdio.h>

int printf(const char * format,
 ...);
```

## **Arguments:**

*format*     A string that controls the format of the output, as described below. The formatting string determines what additional arguments you need to provide.

## **Library:**

`libc`

## **Description:**

The *printf()* function writes output to the *stdout* stream, under control of the argument *format*.



---

If the format string contains invalid multibyte characters, processing stops, and the rest of the format string, including the % characters, is printed. This can happen, for example, if you specify international characters, accents, and diacritical marks using ISO 8859-1 instead of UTF-8. If you call:

```
setlocale(LC_CTYPE, "C-TRADITIONAL");
```

before calling *printf()*, the locale switches multibyte processing from UTF-8 to 1-to-1, and *printf()* safely transfers the misformed multibyte characters.

---

## Format Arguments

If there are leftover arguments after processing *format*, they're ignored.

The *printf()* family of functions allows for language-dependent radix characters. The default character is “.”, but is controlled by `LC_NUMERIC` and *setlocale()*.

## Format control string

The format control string consists of:

### *multibyte characters*

These are copied to the output stream exactly as they occur in the *format* string. An ordinary character in the *format* string is any character, other than a percent character (%), that isn't part of a conversion specifier.

### *conversion specifiers*

These cause argument values to be written as they're encountered during the processing of the *format* string. A conversion specifier is a sequence of characters in the *format* string that begins with “%” and is followed by:

- zero or more *format control flags* that can modify the final effect of the format directive
- an optional decimal integer, or an asterisk (\*), that specifies a *minimum field width* to be reserved for the formatted item
- an optional *precision specification* in the form of a period (.), followed by an optional decimal integer or an asterisk (\*)
- an optional *type length* specification, one of: **h**, **hh**, **j**, **l**, **ll**, **L**, **t** or **z**.
- a character that specifies the type of conversion to be performed. See below.

**Format control flags**

The valid format control flags are:

- Left-justify the formatted item within the field; normally, items are right-justified.
- + Always start a signed, positive object with a plus character (+); normally, only negative items begin with a sign.
- space Always start a signed, positive object with a space character; if both + and a space are specified, + overrides the space.
- # Use an alternate conversion form:
  - For **o** (unsigned octal) conversions, increment the precision, if necessary, so that the first digit is **0**.
  - For **x** or **X** (unsigned hexadecimal) conversions, prepend a nonzero value with **0x** or **0X**.
  - For **e**, **E**, **f**, **g**, or **G** (any floating-point) conversions, always include a decimal-point character in the result, even if no digits follow it; normally, a decimal-point character appears in the result only if there is a digit to follow it.
  - In addition, for **g** or **G** conversions, don't remove trailing zeros from the result.
- 0** (zero) Use leading zeros to pad the field width for **d**, **i**, **o**, **u**, **x**, **X**, **e**, **E**, **f**, **g** and **G** conversions. The “-” flag overrides the this flag.

**Field width**

If you don't specify a field width, or if the given value is less than the number of characters in the converted value (subject to any precision value), a field of sufficient width to contain the converted value is used.

If the converted value has fewer characters than specified by the field width, the value is padded on the left (or right, subject to the



left-justification flag) with spaces or zero characters (0). If the field width begins with a zero, the value is padded with zeros; otherwise, the value is padded with spaces.

If the field width is \* (or 0\*), a value of type `int` from the argument list is used (before a precision argument or a conversion argument) as the minimum field width. A negative field width value is interpreted as a left-justification flag, followed by a positive field width.

**Precision specifier**

As with the field width specifier, a precision specifier of \* causes a value of type `int` from the argument list to be used as the precision specifier. If you give a precision specifier of \*, but there's no precision value in the argument list, a precision of 0 is used.

The precision value affects the following conversions:

- For `d`, `i`, `o`, `u`, `x` and `X` (integer) conversions, the precision specifies the minimum number of digits to appear.
- For `e`, `E` and `f` (fixed-precision, floating-point) conversions, the precision specifies the number of digits to appear after the decimal-point character.
- For `g` and `G` (variable-precision, floating-point) conversions, the precision specifies the maximum number of significant digits to appear.
- For `s` (string) conversions, the precision specifies the maximum number of characters to appear.

**Type length specifier**

A type length specifier affects the conversion as follows:

- `h` causes a `d`, `i`, `o`, `u`, `x` or `X` (integer) format conversion to treat the argument as a `short` or `unsigned short` argument.

Note that, although the argument may have been promoted to an `int` as part of the function call, the value is converted to the smaller type before it's formatted.

- `h` causes an `n` (converted length assignment) operation to assign the converted length to an object of type `short`.

- **hh** is similar to **h**, but treats the argument as a **signed char** or an **unsigned char**.
- **j** causes a **d, i, o, u, x, X** format conversion to process a **intmax\_t** or **uintmax\_t**.
- **j** causes an **n** format conversion to process an **intmax\_t**.
- **L** causes an **a, A, e, E, f, g, G** (double) format conversion to process a **long double** argument.
- **l** (“**el**”) causes a **c** format conversion to process a **wint\_t** argument.
- **l** (“**el**”) causes an **s** format conversion to process a **wchar\_t** argument.
- **l** (“**el**”) causes a **d, i, o, u, x, or X** (integer) format conversion to process a **long** or **unsigned long** argument.
- **l** (“**el**”) causes an **n** (converted length assignment) operation to assign the converted length to an object of type **long**.
- **ll** (double “**el**”) causes a **d, i, o, u, x, or X** (integer) format conversion to assign the converted value to an object of type **long long** or **unsigned long long**.
- **ll** (double “**el**”) causes an **n** (converted length assignment) operation to assign the number of characters that have been read to an object of type **long long**.
- **t** causes a **d, i, o, u, x, X** format conversion to process a **ptrdiff\_t** or the corresponding **unsigned** type argument.
- **t** causes an **n** format conversion to process a **ptrdiff\_t** argument.
- **z** causes a **d, i, o, u, x, X** format conversion to process a **size\_t** argument.
- **z** causes an **n** format conversion to process a **size\_t** argument.

**Conversion  
type  
specifiers**

The valid conversion type specifiers are:

- a, A** Convert an argument of type **double** in the style `[-]0xh.hhhh pld`, where there's one nonzero hexadecimal digit before the decimal point. The number of hexadecimal digits after the decimal point is equal to the precision. If the precision is missing and `FLT_RADIX` is a power of 2, then the precision is sufficient for an exact representation. If the precision is zero and you don't specify the `#` flag, no decimal point is shown.
- The **a** conversion uses the letters **abcdef** and produces **x** and **p**; the **A** conversion **ABCDEF**, **X** and **P**. The exponent always has one digit, even if it's 0, and no more digits than necessary. The values for infinity or NaN are converted in the style of an **f** or **F**.
- c** Convert an argument of type **int** into a value of type **unsigned char** and write the corresponding ASCII character code to the output stream.
- An `l` ("el") qualifier causes a `wint_t` argument to be converted as if by an `ls` conversion into a `wchar_t`, the first element being the `wint_t` and the second being a null wide character.
- d, i** Convert an argument of type **int** into a signed decimal notation and write it to the output stream. The default precision is 1, but if more digits are required, leading zeros are added.
- e, E** Convert an argument of type **double** into a decimal notation in the form `[-]d.ddde [+ | -]dd`. The leading sign appears (subject to the format control flags) only if the argument is negative.
- If the argument is nonzero, the digit before the decimal-point character is nonzero. The precision is used as the number of digits following the decimal-point character. If you don't specify the precision, a default precision of six is used. If the

precision is 0, the decimal-point character is suppressed. The value is rounded to the appropriate number of digits.

The exponent sign and the exponent (that indicates the power of ten by which the decimal fraction is multiplied) are always produced. The exponent is at least two digits long.

For **E** conversions, the exponent begins with the character **E**, rather than **e**.

The arguments infinity or **NaN** are converted in the style of the **f** or **F** conversion specifiers.

- f, F** Convert an argument of type **double** into a decimal notation in the form *[-]ddd.ddd* with the number of digits after the decimal point being equal to the precision specification. The leading sign appears (subject to the format control flags) only if the argument is negative.

The precision is used as the number of digits following the decimal-point character. If you don't specify the precision, a default precision of six is used. If the precision is 0, the decimal-point character is suppressed; otherwise, at least one digit is produced before the decimal-point character. The value is rounded to the appropriate number of digits.

An argument of type **double** that represents infinity or NaN is converted to *[-]inf* or *[-]nan*. The **F** specifier produces *[-]INF* or *[-]NAN*.

- g, G** Convert an argument of type **double** using either the **e** or **f** (or **E**, for a **G** conversion) style of conversion, depending on the value of the argument. In either case, the precision specifies the number of significant digits that are contained in the result. The **e** style conversion is used only if the exponent from such a conversion would be less than -4 or greater than the precision. Trailing zeros are removed from the result, and a decimal-point character only appears if it is followed by a digit.

Arguments representing infinity or NaN are converted in the style of the **f** or **F** conversion specifiers.

- n** Assign the number of characters that have been written to the output stream to the integer pointed to by the argument. No output is produced.
- o** Convert an argument of type **unsigned** into an unsigned octal notation, and write it to the output stream. The default precision is 1, but if more digits are required, leading zeros are added.
- p** Convert an argument of type **void \*** into a value of type **int**, and format the value as for a hexadecimal (**x**) conversion.
- s** Write the characters from the string specified by an argument of type **char \***, up to, but not including the terminating NUL character (`'\0'`), to the output stream. If you specify a precision, no more than that many characters are written.  
  
If you use an **l** (“el”) qualifier, the argument is interpreted as a pointer to a **wchar\_t** array, and each wide character, including the terminating NUL, is converted as if by a call to *wcrtomb()*. The terminating NUL is written only if you don't specify the precision, or if you specify the precision and the length of the character sequence is less than the precision.
- u** Convert an argument of type **unsigned** into an unsigned decimal notation, and write it to the output stream. The default precision is 1, but if more digits are required, leading zeros are added.
- x, X** Convert an argument of type **unsigned** into an unsigned hexadecimal notation, and write it to the output stream. The default precision is 1, but if more digits are required, leading zeros are added.  
  
Hexadecimal notation uses the digits **0** through **9** and the characters **a** through **f** or **A** through **F** for **x** or **X** conversions, respectively, as the hexadecimal digits. Subject to the alternate-form control flag, **0x** or **0X** is prepended to the output.

%        Print a % character (The entire specification is %%).

Any other conversion type specifier character, including another percent character (%), is written to the output stream with no special interpretation.

The arguments must correspond with the conversion type specifiers, left to right in the string; otherwise, indeterminate results will occur.

If the value corresponding to a floating-point specifier is infinity, or not a number (NaN), then the output will be `inf` or `-inf` for infinity, and `nan` or `-nan` for NaNs.

For example, a specifier of the form `%8.*f` defines a field to be at least 8 characters wide, and gets the next argument for the precision to be used in the conversion.

## Returns:

The number of characters written, excluding the terminating NULL, or a negative number if an error occurred (*errno* is set).

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *weekday, *month;

 weekday = "Saturday";
 month = "April";
 printf("%s, %s %d, %d\n", weekday, month, 10, 1999);
 printf("f1 = %8.4f f2 = %10.2E x = %#08x i = %d\n",
 23.45, 3141.5926, 0x1db, -1);
 return EXIT_SUCCESS;
}
```

produces the output:

```
Saturday, April 10, 1999
f1 = 23.4500 f2 = 3.14E+003 x = 0x0001db i = -1
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *fprintf()*, *fwprintf()*, *snprintf()*, *sprintf()*, *swprintf()*, *vfprintf()*, *vfwprintf()*, *vprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*, *wprintf()*

## ***procmgr\_daemon()***

© 2004, QNX Software Systems Ltd.

*Run a process in the background*

### **Synopsis:**

```
#include <sys/procmgr.h>

int procmgr_daemon(int status,
 unsigned flags);
```

### **Arguments:**

- status*    The status that you want to return to the parent process.
- flags*    The flags currently defined (in `<sys/procmgr.h>`) are:
- `PROCMGR_DAEMON_NOCHDIR` — unless this flag is set, *procmgr\_daemon()* changes the current working directory to the root “/”.
  - `PROCMGR_DAEMON_NOCLOSE` — unless this flag is set, *procmgr\_daemon()* closes all file descriptors other than standard input, standard output and standard error.
  - `PROCMGR_DAEMON_NODEVNULL` — unless this flag is set, *procmgr\_daemon()* redirects standard input, standard output and standard error to `/dev/null`.
  - `PROCMGR_DAEMON_KEEPUmask` — unless this flag is set, *procmgr\_daemon()* sets the `umask` to 0 (zero).

### **Library:**

`libc`

### **Description:**

The function *procmgr\_daemon()* function lets programs detach themselves from the controlling terminal and run in the background as system daemons. This also puts the caller into session 1. The argument *status* is returned to the parent process as if *exit()* were called; the returned value is normally `EXIT_SUCCESS`.



**Returns:**

A nonnegative integer, or -1 if an error occurs.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*daemon(), exit(), procmgr\_event\_notify(), procmgr\_event\_trigger(),  
procmgr\_guardian(), procmgr\_session()*

## ***procmgr\_event\_notify()***

© 2004, QNX Software Systems Ltd.

*Ask to be notified of system-wide events*

### **Synopsis:**

```
#include <sys/procmgr.h>

int procmgr_event_notify
 (unsigned flags,
 const struct sigevent * event);
```

### **Arguments:**

*flags*      Flags currently defined in `<sys/procmgr.h>` are:

- `PROC_MGR_EVENT_DAEMON_DEATH` — notify the caller when any process in session 1 dies. This is most useful for watching for the death of daemon processes that use `procmgr_daemon()` to put themselves in session 1 as well as close and redirect file descriptors. As a result of this closing and redirecting, the death of daemons are difficult to detect otherwise.



---

Notification is via the given event, so no information is provided as to which process died. Once you've received the event, you'll need to do something else to find out if processes you care about had died. You can do this by walking through the list of all processes, looking for specific process IDs or process names. If you don't find one, then it has died. The sample code below demonstrates how this can be done.

---

- `PROC_MGR_EVENT_SYNC` — notify the caller of any calls to `sync()` the filesystems.

Setting *flags* to 0 (zero) unarms the event.

*event*      A pointer to a `sigevent` structure that specifies how you want to notified.

### **Library:**

`libc`

**Description:**

The *procmgr\_event\_notify()* function requests that the process manager notify the caller of the system-wide events identified by the given *flags*. A process may have only one notification request active at a time.

**Returns:**

-1 on error; any other value indicates success.

**Examples:**

```
/*
 * This demonstrates procmgr_event_notify() with the
 * PROCMGR_DAEMON_DEATH flag. This flag allows you to
 * be notified if any process in session 1 dies.
 * Daemons are processes that do things that make
 * their death hard to detect (they become daemons by calling
 * procmgr_daemon()). One of the things that happens is that
 * daemons end up in session 1. Hence, the usefulness of the
 * PROCMGR_DAEMON_DEATH flag.
 *
 * When you are notified, you're not told who died.
 * It's up to you to know who should be running. Once notified,
 * you could then walk through the list of which processes are
 * still running and see if all the expected processes are still
 * running. If you know the process id of the processes you
 * are watching out for then this is easiest. If you don't know
 * the process id then your next option may be by process name.
 * The code below does a lookup by process name.
 */

#include <devctl.h>
#include <dirent.h>
#include <errno.h>
#include <fcntl.h>
#include <libgen.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/neutrino.h>
#include <sys/procfs.h>
#include <sys/procmgr.h>

static int check_if_running(char *process);

#define DAEMON_DIED_CODE (_PULSE_CODE_MINAVAIL)
```

```
struct dinfo_s {
 procfs_debuginfo info;
 char pathbuffer[PATH_MAX];
};

int
main(int argc, char **argv)
{
 char *daemon_to_watch;
 int chid, coid, rcvid;
 struct sigevent event;
 struct _pulse msg;

 if (argc != 2) {
 printf("use: %s process_to_watch_for\n", argv[0]);
 exit(EXIT_FAILURE);
 }

 daemon_to_watch = argv[1]; /* the process to watch for */

 chid = ChannelCreate(0);
 coid = ConnectAttach(0, 0, chid, _NTO_SIDE_CHANNEL, 0);
 SIGEV_PULSE_INIT(&event, coid, SIGEV_PULSE_PRIO_INHERIT,
 DAEMON_DIED_CODE, 0);

 /*
 * Ask to be notified via a pulse whenever a
 * daemon process dies
 */

 if (procmgr_event_notify(PROCMGR_EVENT_DAEMON_DEATH,
 &event) == -1) {
 fprintf(stderr, "procmgr_event_notify() failed");
 exit(EXIT_FAILURE);
 }

 while (1) {
 rcvid = MsgReceive(chid, &msg, sizeof(msg), NULL);
 if (rcvid != 0) {
 /* not a pulse, could be unexpected message or error */
 exit(EXIT_FAILURE);
 }

 if (check_if_running(daemon_to_watch) == 0)
 printf("%s is no longer running\n", daemon_to_watch);
 }
 return 0;
}

/*
```

```

* check_if_running - This will walk through all processes
* to see if this particular one is still running.
*/

static int
check_if_running(char *process)
{
 DIR *dirp;
 struct dirent *dire;
 char buffer[20];
 int fd, status;
 pid_t pid;
 struct dinfo_s dinfo;

 if ((dirp = opendir("/proc")) == NULL) {
 perror("Could not open '/proc'");
 return -1;
 }
 while (1) {
 if ((dire = readdir(dirp)) == NULL)
 break;
 if (isdigit(dire->d_name[0])) {
 pid = strtoul(dire->d_name, NULL, 0);

 sprintf(buffer, "/proc/%d/as", pid);
 if ((fd = open(buffer, O_RDONLY)) != NULL) {
 status = devctl(fd, DCMD_PROC_MAPDEBUG_BASE,
 &dinfo, sizeof(dinfo), NULL);
 if (status == EOK) {
 if (!strcmp(process, basename(dinfo.info.path)))
 {
 closedir (dirp);
 return 1;
 }
 }
 /* else some errors are expected, e.g. procnto has
 no MAPDEBUG info and there is a timing issue
 with getting info on the process that died,
 ignore errors */
 close(fd);
 }
 }
 }
 closedir(dirp);
 return 0;
}

```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*procmgr\_daemon()*, *procmgr\_event\_trigger()*, **\_pulse**, **sigevent**

### **Synopsis:**

```
#include <sys/procmgr.h>

int procmgr_event_trigger(unsigned flags);
```

### **Arguments:**

*flags* The type of event that you want to trigger (defined in `<sys/procmgr.h>`):

- `PROC_MGR_EVENT_SYNC` — notify filesystems to `sync()`.

### **Library:**

`libc`

### **Description:**

The function `procmgr_event_trigger()` triggers a system-wide event. The event is sent to all processes that requested (via `procmgr_event_notify()`) to be notified of the event identified by *flags*.

### **Returns:**

-1 on error; any other value indicates success.

### **Examples:**

```
#include <sys/procmgr.h>

int main (void)
{
 procmgr_event_trigger(PROC_MGR_EVENT_SYNC);
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*procmgr\_event\_notify()*, *sync()*



## ***procmgr\_guardian()***

*Let a daemon process take over as a parent*

### **Synopsis:**

```
#include <sys/procmgr.h>

pid_t procmgr_guardian(pid_t pid);
```

### **Arguments:**


*pid* The ID of the child process that you want to become the guardian of the calling process's other children.

### **Library:**

libc

### **Description:**

The function *procmgr\_guardian()* allows a daemon process to declare a child process to take over as parent to its children in the event of its death:



---

*Specifying a guardian for child processes.*

### **Returns:**

-1 on error; any other value on success.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <fcntl.h>
#include <spawn.h>
#include <sys/procmgr.h>
#include <sys/wait.h>

pid_t child = -1;
pid_t guardian = -1;

/*
 * Build a list of the currently running children
 */
void check_children(void) {
 if(child > 0) {
 if(kill(child, 0) == -1) {
 child = -1;
 }
 }
 if(guardian > 0) {
 if(kill(guardian, 0) == -1) {
 guardian = -1;
 }
 }
}

void start_needed_children(void) {
 if(guardian == -1) {
 /* Make a child that will just sit around
 and wait for parent to die */
 while((guardian = fork()) == 0) {
 pid_t parent = getppid();

 /* Wait for parent to die.... */
 fprintf(stderr, "guardian %d waiting on parent %d\n",
 getpid(), parent);

 while(waitpid(parent, 0, 0) != parent);
 /* Then loop around and take over */
 }
 if(guardian == -1) {
 fprintf(stderr, "Unable to start guardian\n");
 } else {
 /* Declare the child a guardian */
 procmgr_guardian(guardian);
 }
 }
}
```

```

 if(child == -1) {
 static char *args[] = { "sleep", "1000000", 0 };

 if((child = spawnp("sleep", 0, 0, 0, args, 0)) == -1) {
 fprintf(stderr, "Couldn't start child\n");
 child = 0; /* don't try again */
 }
 }
 }

int main(int argc, char *argv[]) {
 fprintf(stderr, "parent %d checking children\n", getpid());
 do {
 fprintf(stderr, "checking children\n");

 /* Learn about the newly adopted children */
 check_children();

 /* If anyone is missing, start them */
 start_needed_children();
 } while(wait(0)); /* Then wait for someone to die... */
 return EXIT_SUCCESS;
}

```

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***procmgr\_daemon(), procmgr\_event\_notify(), procmgr\_event\_trigger()*

## ***procmgr\_session()***

© 2004, QNX Software Systems Ltd.

*Provide process manager session support*

### **Synopsis:**

```
#include <sys/procmgr.h>

int procmgr_session(uint32_t nd,
 pid_t sid,
 int id,
 unsigned event);
```

### **Arguments:**

The interpretation of the arguments depends on the *event*.

*nd*        A node descriptor.

*sid*       A session ID.

*id*        A file descriptor, process group, or signal, depending on the event.

*event*    The event; one of:

- PROCMGR\_SESSION\_TCSETSID
- PROCMGR\_SESSION\_SETSID
- PROCMGR\_SESSION\_SETPGRP
- PROCMGR\_SESSION\_SIGNAL\_PID
- PROCMGR\_SESSION\_SIGNAL\_PGRP
- PROCMGR\_SESSION\_SIGNAL\_LEADER

For more information, see below.

### **Library:**

`libc`

## Description:

The *procmgr\_session()* function provides session support to character device terminal drivers in their resource managers, C library functions, and session management applications.

The arguments that you provide need to match the event:

### PROCMGR\_SESSION\_TCSETSID

Used by the *tcsetsid()* function to set the file descriptor, *id*, to be the controlling terminal for the session headed by the session leader, *sid*.

### PROCMGR\_SESSION\_SETSID

Used by the *setsid()* function to create a new session with the calling process becoming the session leader. Pass zero for both *sid* and *id* arguments.

### PROCMGR\_SESSION\_SETPGRP

Used by a character device resource manager to change the process group upon the request of a client calling the *tcsetpgrp()* function. Set the *sid* argument to the client's current session and the *id* argument to the new target process group for the client.

### PROCMGR\_SESSION\_SIGNAL\_PID, PROCMGR\_SESSION\_SIGNAL\_PGRP, PROCMGR\_SESSION\_SIGNAL\_LEADER

Used by a character device resource manager to drop a signal of the type specified as the *id* argument (generally a terminal/job control signal) on the appropriate member of the session specified by the *sid* argument.

## Returns:

- 0 Success.
- 1 Failure.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*setuid(), tcsetpgrp(), tcsetuid()*

## Synopsis:

```
char * __progname
```

## Description:

This global variable holds the basename of the program being executed.



---

This variable isn't defined in any header file. If you want to refer to it, you need to add your own **extern** statement.

---

## Classification:

QNX Neutrino

## See also:

*\_cmdfd()*, *\_cmdname()*

# protoent

© 2004, QNX Software Systems Ltd.

*Structure for information from the protocol database*

## Synopsis:

```
#include <netdb.h>

struct protoent {
 char * p_name;
 char ** p_aliases;
 int p_proto;
};
```

## Description:

The **protoent** structure holds information from the network protocols database, **/etc/protocols**.

The members of this structure are:

|                  |                                                             |
|------------------|-------------------------------------------------------------|
| <i>p_name</i>    | The name of the protocol.                                   |
| <i>p_aliases</i> | A zero-terminated list of alternate names for the protocol. |
| <i>p_proto</i>   | The protocol number.                                        |

## Classification:

Unix, POSIX 1003.1-2001

## See also:

*endprotoent()*, *getprotobyname()*, *getprotobynumber()*, *getprotoent()*, *setprotoent()*

**/etc/protocols** in the *Utilities Reference*



### **Synopsis:**

```
#include <pthread.h>

int pthread_abort(pthread_t thread);
```

### **Arguments:**

*thread*     The ID of the thread that you want to terminate, which you can get when you call *pthread\_create()* or *pthread\_self()*.

### **Library:**

`libc`

### **Description:**

The *pthread\_abort()* function terminates the target thread. Termination takes effect immediately and isn't a function of the cancelability state of the target thread. No cancellation handlers or thread-specific-data destructor functions are executed. Thread abortion doesn't release any application-visible process resources, including, but not limited to, mutexes and file descriptors. (The behavior of POSIX calls following a call to *pthread\_abort()* is unspecified.)

The status of PTHREAD\_ABORTED is available to any thread joining with the target thread. The constant PTHREAD\_ABORTED expands to a constant expression, of type `void *`. Its value doesn't match any pointer to an object in memory, or the values NULL and PTHREAD\_CANCELED.

The side effects of aborting a thread that's suspended during a call of a POSIX 1003.1 function are the same as the side effects that may be seen in a single-threaded process when a call to a POSIX 1003.1 function is interrupted by a signal and the given function returns EINTR. Any such side effects occur before the thread terminates.

**Returns:**

EOK        Success.  
ESRCH      No thread with the given ID was found.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cancel()*, *pthread\_detach()*, *pthread\_exit()*, *ThreadDestroy()*

### **Synopsis:**

```
#include <process.h>

int pthread_atfork(void (*prepare) (void),
 void (*parent) (void),
 void (*child) (void));
```

### **Arguments:**

*prepare*      NULL, or a pointer to the handler to call before the fork.

*parent*        NULL, or a pointer to the handler to call after the fork in the parent process.

*child*         NULL, or a pointer to the handler to call after the fork in the child process.

### **Library:**

`libc`

### **Description:**

The *pthread\_atfork()* function registers fork handler functions to be called before and after a *fork()*, in the context of the thread that called *fork()*. You can set one or more of the arguments to NULL to indicate no handler.

You can register multiple *prepare*, *parent*, and *child* fork handlers, by making additional calls to *pthread\_atfork()*. In this case, the *parent* and *child* handlers are called in the order they were registered, and the *prepare* handlers are called in the reverse order.



---

You can't use the *pthread\_atfork()* function for useful purposes as the C library doesn't have the necessary handlers. It also implies that Neutrino currently doesn't support *fork()* in multi-threaded programs.

---

**Returns:**

EOK            Success.  
ENOMEM        Insufficient memory to record fork handlers.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atexit()*, *fork()*

## ***pthread\_attr\_destroy()***

*Destroy a thread-attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_destroy(pthread_attr_t * attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_attr_t` structure that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_destroy()` function destroys the given thread-attribute object.



---

The QNX implementation of this function doesn't actually free the memory used by the `pthread_attr_t` structure. To conform to the POSIX standard, don't reuse the attribute object unless you reinitialize it by calling `pthread_attr_init()`.

---

You can use a thread-attribute object to define the attributes of new threads when you call `pthread_create()`.

### **Returns:**

0 for success, or an error number.

### **Errors:**

EOK Success.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_init()*, *pthread\_create()*

## ***pthread\_attr\_getdetachstate()***

*Get thread detach state attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getdetachstate(
 const pthread_attr_t* attr,
 int* detachstate);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*detachstate*    A pointer to a location where the function can store the thread detach state. For more information, see `pthread_attr_setdetachstate()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getdetachstate()` function gets the thread detach state attribute from the thread attribute object *attr* and returns it in *detachstate*.

### **Returns:**

EOK    Success.

### **Classification:**

POSIX 1003.1 (Threads)

## *pthread\_attr\_getdetachstate()*

© 2004, QNX Software Systems Ltd.

---

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **See also:**

*pthread\_attr\_setdetachstate()*, *pthread\_attr\_init()*, *pthread\_create()*.



## ***pthread\_attr\_getguardsize()***

*Get the size of the thread's guard area*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getguardsize(
 const pthread_attr_t* attr,
 size_t* guardsize);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*guardsize*      A pointer to a location where the function can store the size of the thread's guard area. For more information, see `pthread_attr_setguardsize()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getguardsize()` function gets the value of the thread `guardsize` attribute from the attribute structure `attr`.

### **Returns:**

EOK            Success.

EINVAL        Invalid pointer, `attr`, to a `pthread_attr_t` structure.

### **Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_setguardsize()*

## ***pthread\_attr\_getinheritsched()***

*Get a thread's inherit-scheduling attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getinheritsched(
 const pthread_attr_t* attr,
 int* inheritsched);
```

### **Arguments:**

*attr*                    A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*inheritsched*           A pointer to a location where the function can store the value of the inherit-scheduling attribute. For more information, see `pthread_attr_setinheritsched()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getinheritsched()` function gets the thread inherit-scheduling attribute from the attribute object *attr* and returns it in *inheritsched*.

The inherit-scheduling attribute determines whether a thread inherits the scheduling policy of its parent.

### **Returns:**

EOK      Success.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_setinheritsched()*, *pthread\_attr\_init()*, *pthread\_create()*

## ***pthread\_attr\_getschedparam()***

*Get thread scheduling parameters attribute*

### **Synopsis:**

```
#include <pthread.h>
#include <sched.h>

int pthread_attr_getschedparam(
 const pthread_attr_t * attr,
 struct sched_param * param);
```

### **Arguments:**

*attr*      A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*param*     A pointer to a `sched_param` structure where the function can store the current scheduling parameters.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getschedparam()` function gets the thread scheduling parameters attribute from the thread attribute object *attr* and returns it in *param*.

### **Returns:**

EOK      Success.

### **Classification:**

POSIX 1003.1 (Threads)

## *pthread\_attr\_getschedparam()*

© 2004, QNX Software Systems Ltd.

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **See also:**

*pthread\_attr\_setschedparam()*, *pthread\_attr\_init()*, *pthread\_create()*,  
**sched\_param**

### **Synopsis:**

```
#include <pthread.h>
#include <sched.h>

int pthread_attr_getschedpolicy(
 const pthread_attr_t* attr,
 int* policy);
```

### **Arguments:**

*attr* A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*policy* The current thread scheduling policy. For more information, see `pthread_attr_setschedpolicy()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getschedpolicy()` function gets the thread scheduling policy attribute from the thread attribute object *attr* and returns it in *policy*.

### **Returns:**

EOK Success.

### **Classification:**

POSIX 1003.1 (Threads)

## *pthread\_attr\_getschedpolicy()*

© 2004, QNX Software Systems Ltd.

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **See also:**

*pthread\_attr\_setschedpolicy()*, *pthread\_attr\_init()*, *pthread\_create()*.



### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getscope(
 const pthread_attr_t *attr,
 int *scope);
```

### **Arguments:**

*attr* A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*scope* A pointer to a location where the function can store the current contention scope. For more information, see `pthread_attr_setscope()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getscope()` function gets the thread contention scope attribute from the thread attribute object *attr* and returns it in *scope*.

### **Returns:**

EOK Success.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_attr\_setscope()*, *pthread\_attr\_init()*, *pthread\_create()*.

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getstackaddr(
 const pthread_attr_t* attr,
 void** stackaddr);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*stackaddr*      A pointer to a location where the function can store the address of the thread stack.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getstackaddr()` function gets the thread stack address attribute from the thread attribute object *attr* and returns it in *stackaddr*.

For more information about the thread stack, see `pthread_attr_setstackaddr()`

### **Returns:**

EOK      Success.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_setstackaddr()*, *pthread\_attr\_init()*, *pthread\_create()*.

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getstacklazy(
 const pthread_attr_t * attr,
 int *lazystack);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*lazystack*      A pointer to a location where the function can store the current lazy-stack attribute. For more information, see `pthread_attr_setstacklazy()`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getstacklazy()` function gets the thread lazy-stack attribute in the attribute object *attr* and stores it in the location pointed to by *lazystack*.

### **Returns:**

EOK      Success.

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_setinheritsched()*, *pthread\_attr\_setstacklazy()*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_getstacksize(
 const pthread_attr_t* attr,
 size_t* stacksize);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*stacksize*      A pointer to a location where the function can store the stack size to be used for new threads.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_getstacksize()` function gets the thread stack size attribute from the thread attribute object *attr* and returns it in *stacksize*.

You can set the stack size by calling `pthread_attr_setstacksize()`.

### **Returns:**

EOK      Success.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_setstacksize(), pthread\_attr\_init(), pthread\_create()*.



## Synopsis:

```
#include <pthread.h>

int pthread_attr_init(pthread_attr_t *attr);
```

## Arguments:

*attr* A pointer to the `pthread_attr_t` structure that you want to initialize. For more information, see below.

## Library:

`libc`

## Description:

The `pthread_attr_init()` function initializes the thread attributes in the thread attribute object *attr* to their default values:

| <b>Attribute</b>       | <b>Default Value</b>         |
|------------------------|------------------------------|
| <i>detachstate</i>     | PTHREAD_CREATE_JOINABLE      |
| <i>schedpolicy</i>     | PTHREAD_INHERIT_SCHED        |
| <i>schedparam</i>      | Inherited from parent thread |
| <i>contentionscope</i> | PTHREAD_SCOPE_SYSTEM         |
| <i>stacksize</i>       | 4K bytes                     |
| <i>stackaddr</i>       | NULL                         |

After initialization, you can use the `pthread_attr_*` family of functions to get and set the attributes:

| <b>Get</b>                            | <b>Set</b>                            |
|---------------------------------------|---------------------------------------|
| <i>pthread_attr_getdetachstate()</i>  | <i>pthread_attr_setdetachstate()</i>  |
| <i>pthread_attr_getguardsize()</i>    | <i>pthread_attr_setguardsize()</i>    |
| <i>pthread_attr_getinheritsched()</i> | <i>pthread_attr_setinheritsched()</i> |
| <i>pthread_attr_getschedparam()</i>   | <i>pthread_attr_setschedparam()</i>   |
| <i>pthread_attr_getschedpolicy()</i>  | <i>pthread_attr_setschedpolicy()</i>  |
| <i>pthread_attr_getscope()</i>        | <i>pthread_attr_setscope()</i>        |
| <i>pthread_attr_getstackaddr()</i>    | <i>pthread_attr_setstackaddr()</i>    |
| <i>pthread_attr_getstacklazy()</i>    | <i>pthread_attr_setstacklazy()</i>    |
| <i>pthread_attr_getstacksize()</i>    | <i>pthread_attr_setstacksize()</i>    |

You can also set some non-POSIX attributes; for more information, see “QNX extensions,” in the documentation for *pthread\_create()*.

You can then pass the attribute object to *pthread\_create()* to create a thread with the required attributes. You can use the same attribute object in multiple calls to *pthread\_create()*.

The effect of initializing an already-initialized thread-attribute object is undefined.

**Returns:**

EOK     Success.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*pthread\_attr\_destroy()*, *pthread\_create()*

## ***pthread\_attr\_setdetachstate()***

© 2004, QNX Software Systems Ltd.

*Set thread detach state attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setdetachstate(
 pthread_attr_t* attr,
 int detachstate);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*detachstate*    The new value for the thread detach state:

- `PTHREAD_CREATE_JOINABLE` — create the thread in a joinable state.
- `PTHREAD_CREATE_DETACHED` — create the thread in a detached state.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setdetachstate()` function sets the thread detach state attribute in the thread attribute object *attr* to *detachstate*.

The default value for the thread detach state is `PTHREAD_CREATE_JOINABLE`.

### **Returns:**

`EOK`            Success.

`EINVAL`        Invalid thread detach state value.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_getdetachstate()*, *pthread\_attr\_init()*, *pthread\_create()*,  
*pthread\_detach()*, *pthread\_join()*.

## ***pthread\_attr\_setguardsize()***

© 2004, QNX Software Systems Ltd.

*Set the size of the thread's guard area*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setguardsize(
 pthread_attr_t* attr,
 size_t guardsize);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*guardsize*      The new value for the size of the thread's guard area.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setguardsize()` function sets the size of the thread's guard area in the attribute structure *attr* to *guardsize*.

If *guardsize* is 0, threads created with *attr* have no guard area; otherwise, a guard area of at least *guardsize* bytes is provided. You can get the default *guardsize* value by specifying `_SC_PAGESIZE` in a call to `sysconf()`.

The *guardsize* attribute controls the size of the guard area for the thread's stack. This guard area helps protect against stack overflows; *guardsize* bytes of extra memory is allocated at the overflow end of the stack. If a thread overflows into this buffer, it receives a SIGSEGV signal.

The *guardsize* attribute is provided because:

- Stack overflow protection can waste system resources. An application that creates many threads can save system resources by

turning off guard areas if it trusts its threads not to overflow the stack.

- When threads allocate large objects on the stack, a large *guardsize* is required to detect stack overflows.

### Returns:

|        |                                                                                                            |
|--------|------------------------------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                                   |
| EINVAL | Invalid pointer, <i>attr</i> , to a <code>pthread_attr_t</code> structure, or <i>guardsize</i> is invalid. |

### Classification:

Standard Unix

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### Caveats:

If you provide a stack (using *attr*'s *stackaddr* attribute; see *pthread\_attr\_setstackaddr()*), the *guardsize* is ignored, and there's no stack overflow protection for that thread.

The *guardsize* argument is completely ignored when using a physical mode memory manager.

## *pthread\_attr\_setguardsize()*

---

© 2004, QNX Software Systems Ltd.

### **See also:**

*pthread\_attr\_getguardsize()*, *pthread\_attr\_init()*,  
*pthread\_attr\_setstackaddr()*, *sysconf()*



## ***pthread\_attr\_setinheritsched()***

*Set a thread's inherit-scheduling attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setinheritsched(
 pthread_attr_t * attr,
 int inheritsched);
```

### **Arguments:**

- |                     |                                                                                                                                                                                                                                                                                                                                        |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>attr</i>         | A pointer to the <code>pthread_attr_t</code> structure that defines the attributes to use when creating new threads. For more information, see <code>pthread_attr_init()</code> .                                                                                                                                                      |
| <i>inheritsched</i> | The new value for the thread's inherit-scheduling attribute: <ul style="list-style-type: none"><li>• <code>PTHREAD_INHERIT_SCHED</code> — the thread inherits the scheduling policy of the parent thread.</li><li>• <code>PTHREAD_EXPLICIT_SCHED</code> — use the scheduling policy specified in <i>attr</i> for the thread.</li></ul> |

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setinheritsched()` function sets the thread inherit scheduling attribute in the attribute object *attr* to *inheritsched*.

The default value of the thread inherit scheduling attribute is `PTHREAD_INHERIT_SCHED`.

**Returns:**

- EOK            Success.
- EINVAL        Invalid thread attribute object *attr*.
- ENOTSUP      Invalid thread inherit scheduling attribute *inheritsched*.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*pthread\_attr\_getinheritsched()*, *pthread\_attr\_init()*, *pthread\_create()*.

## ***pthread\_attr\_setschedparam()***

*Set a thread's scheduling parameters attribute*

### **Synopsis:**

```
#include <pthread.h>
#include <sched.h>

int pthread_attr_setschedparam(
 pthread_attr_t * attr,
 const struct sched_param * param);
```

### **Arguments:**

*attr* A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*param* A pointer to a `struct sched_param` structure that defines the thread's scheduling parameters.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setschedparam()` function sets the thread scheduling parameters attribute in the thread attribute object *attr* to *param*.

The thread scheduling parameters are used only if you've set the thread inherit scheduling attribute to `PTHREAD_EXPLICIT_SCHED` by calling `pthread_attr_setinheritsched()`. By default, a thread inherits its parent's priority.

### **Returns:**

EOK Success.

EINVAL Invalid thread attribute object *attr*.

ENOTSUP Invalid thread scheduling parameters attribute *param*.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_getschedparam()*, *pthread\_attr\_setinheritsched()*,  
*pthread\_attr\_init()*, *pthread\_create()*, **sched\_param**

**Synopsis:**

```
#include <pthread.h>
#include <sched.h>

int pthread_attr_setschedpolicy(
 pthread_attr_t* attr,
 int policy);
```

**Arguments:**

*attr* A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*policy* The new value for the scheduling policy:

- SCHED\_FIFO — first-in first-out scheduling.
- SCHED\_RR — round-robin scheduling.
- SCHED\_OTHER — currently the same as SCHED\_RR.
- SCHED\_NOCHANGE — don't change the policy.
- SCHED\_SPORADIC — sporadic scheduling.

**Library:**

`libc`

**Description:**

The `pthread_attr_setschedpolicy()` function sets the thread scheduling policy attribute in the thread attribute object *attr* to *policy*.

The *policy* attribute is used only if you've set the thread inherit-scheduling attribute to `PTHREAD_EXPLICIT_SCHED` by calling `pthread_attr_setinheritsched()`.

For descriptions of the scheduling policies, see "Scheduling algorithms" in the chapter on the Neutrino microkernel in the *System Architecture* guide.

**Returns:**

|         |                                                  |
|---------|--------------------------------------------------|
| EOK     | Success.                                         |
| EINVAL  | Invalid thread attribute object <i>attr</i> .    |
| ENOTSUP | Invalid thread scheduling policy <i>policy</i> . |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_getschedpolicy()*, *pthread\_attr\_init()*, *pthread\_create()*.

## ***pthread\_attr\_setscope()***

*Set thread contention scope attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setscope(pthread_attr_t* attr,
 int scope);
```

### **Arguments:**

- attr*      A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.
- scope*     The new value for the contention scope attribute:
- `PTHREAD_SCOPE_SYSTEM` — schedule all threads together.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setscope()` sets the thread contention scope attribute in the thread attribute object *attr* to *scope*.

### **Returns:**

- `EOK`            Success.
- `EINVAL`        Invalid thread attribute object *attr*.
- `ENOTSUP`      Invalid thread contention scope attribute *scope*.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_getscope(), pthread\_attr\_init(), pthread\_create()*.



## ***pthread\_attr\_setstackaddr()***

*Set the thread stack address attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setstackaddr(pthread_attr_t * attr,
 void * stackaddr);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see *pthread\_attr\_init()*.

*stackaddr*      A pointer to the block of memory that you want a new thread to use as its stack.

### **Library:**

`libc`

### **Description:**

The *pthread\_attr\_setstackaddr()* function sets the thread stack address attribute in the attribute object *attr* to *stackaddr*.

The default value for the thread stack address attribute is NULL. A thread created with a NULL stack address attribute will have a stack dynamically allocated by the system of minimum size `PTHREAD_STACK_MIN`. If the system allocates a stack, it reclaims the space when the thread terminates. If you allocate a stack, you must free it.

### **Returns:**

EOK      Success.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The QNX interpretation of PTHREAD\_STACK\_MIN is enough memory to run a thread that does nothing:

```
void nothingthread(void)
{
 return;
}
```

**See also:**

*pthread\_attr\_getstackaddr(), pthread\_attr\_init(), pthread\_create().*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setstacklazy(
 pthread_attr_t * attr,
 int lazystack);
```

### **Arguments:**

- attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.
- lazystack*       One of:
- `PTHREAD_STACK_LAZY` — allocate physical memory for the thread stack on demand (the default).
  - `PTHREAD_STACK_NOTLAZY` — allocate physical memory for the whole stack up front. Use this value to ensure that your server processes don't die later on because they're unable to allocate stack memory. We recommend that you set the stack size as well, because the default stack size is probably much larger than you really need.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setstacklazy()` function sets the thread stack attribute in the attribute object *attr* to *lazystack*.

**Returns:**

- EOK Success.
- EINVAL The thread-attribute object that *attr* points to is invalid.
- ENOTSUP The value of *lazystack* is invalid.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point No
- Interrupt handler No
- Signal handler Yes
- Thread Yes

**See also:**

*pthread\_attr\_getstacklazy()*, *pthread\_attr\_setinheritsched()*

## ***pthread\_attr\_setstacksize()***

*Set the thread stack-size attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_attr_setstacksize(pthread_attr_t * attr,
 size_t stacksize);
```

### **Arguments:**

*attr*            A pointer to the `pthread_attr_t` structure that defines the attributes to use when creating new threads. For more information, see `pthread_attr_init()`.

*stacksize*      The size of the stack you want to use in new threads. The minimum value of the thread stack-size attribute is `PTHREAD_STACK_MIN`.

### **Library:**

`libc`

### **Description:**

The `pthread_attr_setstacksize()` function sets the thread stack size attribute in the thread attribute object *attr* to *stacksize*.

### **Returns:**

EOK            Success.

EINVAL        The value of *stacksize* is less than `PTHREAD_STACK_MIN` or greater than the system limit.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The QNX interpretation of PTHREAD\_STACK\_MIN is enough memory to run a thread that does nothing:

```
void nothingthread(void)
{
 return;
}
```

**See also:**

*pthread\_attr\_getstacksize()*, *pthread\_attr\_init()*, *pthread\_create()*.

## ***pthread\_barrier\_destroy()***

*Destroy a barrier object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_barrier_destroy(pthread_barrier_t * barrier);
```

### **Arguments:**

*barrier*     A pointer to the `pthread_barrier_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_barrier_destroy()` function destroys the barrier referenced by *barrier* and releases any resources used by the barrier. Subsequent use of the barrier is undefined until you reinitialize the barrier by calling `pthread_barrier_init()`.

### **Returns:**

`EBUSY`     The *barrier* is in use.  
`EINVAL`     Invalid *barrier*.  
`EOK`        Success.

### **Classification:**

POSIX 1003.1j (draft)

#### **Safety**

---

Cancellation point    No  
Interrupt handler      No

*continued...*

## *pthread\_barrier\_destroy()*

© 2004, QNX Software Systems Ltd.

---

### **Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

### **See also:**

*pthread\_barrier\_init()*, *pthread\_barrier\_wait()*



**Synopsis:**

```
#include <pthread.h>

int pthread_barrier_init(
 pthread_barrier_t * barrier,
 const pthread_barrierattr_t * attr
 unsigned int count);
```

**Arguments:**

- barrier* A pointer to the `pthread_barrier_t` object that you want to initialize.
- attr* NULL, or a pointer to a `pthread_barrierattr_t` structure that specifies the attributes that you want to use for the barrier.
- count* The number of threads that must call `pthread_barrier_wait()` before any of them successfully returns from the call. This value must be greater than zero.

**Library:**

`libc`

**Description:**

The `pthread_barrier_init()` function allocates any resources required to use the barrier referenced by *barrier* and initializes the barrier with attributes referenced by *attr*. If *attr* is NULL, the default barrier attributes are used. The effect is the same as passing the address of a default barrier attributes object. Once it's initialized, you can use the barrier any number of times without reinitializing it.

If `pthread_barrier_init()` fails, the barrier isn't initialized.

In cases where the default barrier attributes are appropriate, you can use `PTHREAD_BARRIER_INITIALIZER()` macro to initialize barriers

that are statically allocated. The effect is equivalent to dynamic initialization by a call to *pthread\_barrier\_init()* with parameter *attr* specified as NULL, except that no error checks are performed.

**Returns:**

- EAGAIN      The system lacks the necessary resources to initialize another barrier.
- EBUSY      Attempt to reinitialize a barrier while it's in use.
- EFAULT      A fault occurred when the kernel tried to access *barrier* or *attr*.
- EINVAL      Invalid value specified by *attr*.
- EOK         Success.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                Yes

**See also:**

*pthread\_barrierattr\_init()*, *pthread\_barrier\_destroy()*,  
*pthread\_barrier\_wait()*

## ***pthread\_barrier\_wait()***

*Synchronize participating threads at the barrier*

### **Synopsis:**

```
#include <sync.h>

int pthread_barrier_wait(pthread_barrier_t * barrier);
```

### **Arguments:**

*barrier* A pointer to the `pthread_barrier_t` object that you want to use to synchronize the threads. You must initialize the barrier by calling `pthread_barrier_init()`, before calling `pthread_barrier_wait()`.

### **Library:**

`libc`

### **Description:**

The `pthread_barrier_wait()` function synchronizes participating threads at the barrier referenced by *barrier*. The calling thread blocks — that is, doesn't return from `pthread_barrier_wait()` — until the required number of threads have called `pthread_barrier_wait()`, specifying the barrier.

When the required number of threads have called `pthread_barrier_wait()` specifying the barrier, the constant `BARRIER_SERIAL_THREAD` is returned to one unspecified thread, and zero is returned to each of the remaining threads. At this point, the barrier is reset to the state it occupied as a result of the most recent `pthread_barrier_init()` function that referenced it.

The constant `BARRIER_SERIAL_THREAD` is defined in `<pthread.h>`, and its value is distinct from any other value that `pthread_barrier_wait()` returns.

If a signal is delivered to a thread blocked on a barrier, on return from the signal handler, the thread resumes waiting at the barrier as if it hadn't been interrupted.

**Returns:**

BARRIER\_SERIAL\_THREAD for a single (arbitrary) thread synchronized at the barrier and zero for each of the other threads; otherwise, an error number is returned:

EINVAL      The *barrier* argument isn't initialized.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_barrier\_destroy()*, *pthread\_barrier\_init()*

## ***pthread\_barrierattr\_destroy()***

*Destroy a barrier-attributes object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_barrierattr_destroy(
 pthread_barrierattr_t * attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_barrierattr_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_barrierattr_destroy()` function destroys the barrier-attributes object *attr*. Subsequent use of the object is undefined until you reinitialize the object by calling `pthread_barrierattr_init()`.

Once you've used a barrier-attributes object to initialize one or more barriers, any changes to the attributes object (including destroying it) don't affect any previously initialized barriers.

### **Returns:**

EOK Success.  
EINVAL Invalid barrier attribute object *attr*.

### **Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_barrier\_destroy()*, *pthread\_barrierattr\_getpshared()*,  
*pthread\_barrierattr\_init()*, *pthread\_barrierattr\_setpshared()*

## ***pthread\_barrierattr\_getpshared()***

*Get the value of a barrier's process-shared attribute*

### **Synopsis:**

```
#include <pthread.h>

int pthread_barrierattr_getpshared(
 const pthread_barrierattr_t * attr
 int * pshared);
```

### **Arguments:**

*attr*            A pointer to the `pthread_barrierattr_t` object whose attribute you want to query. You must have initialized this object by calling `pthread_barrierattr_init()` before calling `pthread_barrierattr_getpshared()`.

*pshared*        A pointer to a location where the function can store the value of the process-shared attribute. For information about the possible values, see `pthread_barrierattr_setpshared()`.

### **Library:**

`libc`

### **Description:**

The `pthread_barrierattr_getpshared()` function gets the value of the process-shared attribute from the attributes object referenced by *attr* and stores the value in the object referenced by *pshared*.

### **Returns:**

EOK            Success.

EINVAL        Invalid barrier attribute object *attr*.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_barrier\_init()*, *pthread\_barrierattr\_destroy()*,  
*pthread\_barrierattr\_init()*, *pthread\_barrierattr\_setpshared()*



**Synopsis:**

```
#include <pthread.h>

int pthread_barrierattr_init(
 pthread_barrierattr_t * attr);
```

**Arguments:**

*attr* A pointer to the `pthread_barrierattr_t` object that you want to initialize.

**Library:**

`libc`

**Description:**

The `pthread_barrierattr_init()` function initializes the barrier attributes object *attr* with the default value for all of the attributes.

**Returns:**

|        |                                                                          |
|--------|--------------------------------------------------------------------------|
| EOK    | Success.                                                                 |
| ENOMEM | Insufficient memory to initialize barrier attribute object <i>attr</i> . |

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_barrier\_init()*, *pthread\_barrierattr\_destroy()*,  
*pthread\_barrierattr\_getpshared()*, *pthread\_barrierattr\_setpshared()*

## ***pthread\_barrierattr\_setpshared()***

*Set the value of a barrier's process-shared attribute*

### **Synopsis:**

```
#include <sync.h>

int pthread_barrierattr_setpshared(
 pthread_barrierattr_t * attr
 int pshared);
```

### **Arguments:**

- attr* A pointer to the `pthread_barrierattr_t` object whose attribute you want to set. You must have initialized this object by calling `pthread_barrierattr_init()` before calling `pthread_barrierattr_setpshared()`.
- pshared* The new value of the process-shared attribute; one of:
- `PTHREAD_PROCESS_SHARED` — allow a barrier to be operated upon by any thread that has access to the memory where the barrier is allocated.
  - `PTHREAD_PROCESS_PRIVATE` (default behavior) — allow a barrier to be operated on only by threads created within the same process as the thread that initialized the barrier. If threads of different processes attempt to operate on such a barrier, the behavior is as if `PTHREAD_PROCESS_SHARED` were set.

### **Library:**

`libc`

### **Description:**

The `pthread_barrierattr_setpshared()` function sets the process-shared attribute in an initialized attributes object referenced by *attr*.

**Returns:**

- EOK Success.
- EINVAL The attribute object, *attr*, or the new value specified in *pshared* isn't valid.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

- Cancellation point No
- Interrupt handler No
- Signal handler Yes
- Thread Yes

**See also:**

*pthread\_barrier\_init()* *pthread\_barrierattr\_destroy()*,  
*pthread\_barrierattr\_getpshared()*, *pthread\_barrierattr\_init()*

## Synopsis:

```
#include <pthread.h>

int pthread_cancel(pthread_t thread);
```

## Arguments:

*thread*      The ID of the thread that you want to cancel, which you can get when you call *pthread\_create()* or *pthread\_self()*.

## Library:

`libc`

## Description:

The *pthread\_cancel()* function requests that the target thread *thread* be canceled (terminated). The cancellation type and state of the target determine when the cancellation takes effect.

When the cancellation is acted on, the target's cancellation cleanup handlers are called. When the last cancellation cleanup handler returns, the target's thread-specific-data destructor functions are called. When the last destructor function returns, the target is terminated. Cancellation processing in the target thread runs asynchronously with respect to the calling thread.

## Returns:

EOK          Success.

ESRCH        No thread with thread ID *thread* exists.

## Classification:

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cleanup\_push()*, *pthread\_cleanup\_pop()*, *pthread\_cond\_wait()*,  
*pthread\_cond\_timedwait()*, *pthread\_exit()*, *pthread\_join()*,  
*pthread\_key\_create()*, *pthread\_setcancelstate()*,  
*pthread\_setcanceltype()*, *pthread\_testcancel()*, *ThreadCancel()*.

### **Synopsis:**

```
#include <pthread.h>

void pthread_cleanup_pop(int execute);
```

### **Arguments:**

*execute*      Zero if you don't want to execute the handler; nonzero if you do.

### **Library:**

`libc`

### **Description:**

The *pthread\_cleanup\_pop()* macro pops the top cancellation-cleanup handler from the calling thread's cancellation-cleanup stack and invokes the handler if *execute* is nonzero.



---

The *pthread\_cleanup\_pop()* macro expands to a few lines of code that end with a closing brace (`}`), but don't have a matching opening brace (`{`). You must pair *pthread\_cleanup\_pop()* with *pthread\_cleanup\_push()* within the same lexical scope.

---

### **Examples:**

See *pthread\_cleanup\_push()*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_cleanup\_push()*, *pthread\_cancel()*, *pthread\_exit()*



### **Synopsis:**

```
#include <pthread.h>

void pthread_cleanup_push(void (routine) (void*),
 void* arg);
```

### **Arguments:**

*routine*    The handler that you want to push onto the thread's stack.

*arg*        A pointer to whatever data you want to pass to the function when it's called.

### **Library:**

`libc`

### **Description:**

The *pthread\_cleanup\_push()* function pushes the given cancellation-cleanup handler *routine* onto the calling thread's cancellation-cleanup stack.

The cancellation-cleanup handler is popped from the stack and invoked with argument *arg* when the thread:

- exits (i.e. calls *pthread\_exit()*)
- acts on a cancellation request
- calls *pthread\_cleanup\_pop()* with a nonzero argument.



---

The *pthread\_cleanup\_push()* macro expands to a few lines of code that start with an opening brace (`{`), but don't have a matching closing brace (`}`). You must pair *pthread\_cleanup\_push()* with *pthread\_cleanup\_pop()* within the same lexical scope.

---

## Examples:

Use a cancellation cleanup handler to free resources, such as a mutex, when a thread is terminated:

```
#include <pthread.h>

pthread_mutex_t lock = PTHREAD_MUTEX_INITIALIZER;

void unlock(void * arg)
{
 pthread_mutex_unlock(&lock);
}

void * function(void * arg)
{
 while(1)
 {
 pthread_mutex_lock(&lock);
 pthread_cleanup_push(&unlock, NULL);

 /*
 * Any of the possible cancellation points could
 * go here.
 */
 pthread_testcancel();

 pthread_cleanup_pop(1);
 }
}
```

## Classification:

POSIX 1003.1 (Threads)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cleanup\_pop()*, *pthread\_cancel()*, *pthread\_exit()*

## ***pthread\_cond\_broadcast()***

© 2004, QNX Software Systems Ltd.

*Unblock threads waiting on condition*

### **Synopsis:**

```
#include <pthread.h>

int pthread_cond_broadcast(pthread_cond_t* cond);
```

### **Arguments:**

*cond* A pointer to the `pthread_cond_t` object for which you want to unblock the threads.

### **Library:**

`libc`

### **Description:**

The `pthread_cond_broadcast()` function unblocks all threads currently blocked on the condition variable *cond*. The threads are unblocked in priority order.

If more than one thread at a particular priority is blocked, those threads are unblocked in FIFO order.

### **Returns:**

EOK Success.

EFAULT A fault occurred trying to access the buffers provided.

EINVAL Invalid condition variable *cond*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_cond\_signal()*, *pthread\_cond\_wait()*, *SyncCondvarSignal()*

## ***pthread\_cond\_destroy()***

© 2004, QNX Software Systems Ltd.

*Destroy condition variable*

### **Synopsis:**

```
#include <pthread.h>

int pthread_cond_destroy(pthread_cond_t* cond);
```

### **Arguments:**

*cond* A pointer to the `pthread_cond_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_cond_destroy()` function destroys the condition variable *cond*. After you've destroyed a condition variable, you shouldn't reuse it until you've reinitialized it by calling `pthread_cond_init()`.

### **Returns:**

EOK Success.

EBUSY Another thread is blocked on the condition variable *cond*.

EINVAL Invalid condition variable *cond*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*pthread\_cond\_init()*, *SyncDestroy()*

## ***pthread\_cond\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize a condition variable*

### **Synopsis:**

```
#include <pthread.h>

pthread_cond_t cond = PTHREAD_COND_INITIALIZER;

int pthread_cond_init(pthread_cond_t* cond,
 pthread_condattr_t* attr);
```

### **Arguments:**

*cond* A pointer to the `pthread_cond_t` object that you want to initialize.

*attr* NULL, or a pointer to a `pthread_condattr_t` object that specifies the attributes that you want to use for the condvar. For more information, see `pthread_condattr_init()`.

### **Library:**

`libc`

### **Description:**

The `pthread_cond_init()` function initializes the condition variable `cond` with the attributes in the condition variable attribute object `attr`. If `attr` is NULL, `cond` is initialized with the default values for the attributes.

If the condition variable is statically allocated, you can initialize it with the default attribute values by assigning to it the macro `PTHREAD_COND_INITIALIZER`.

Condition variables have at least the following attributes defined:

`PTHREAD_PROCESS_PRIVATE`

The condition variable can only be accessed by threads created within the same process as the thread that initialized the condition variable.



**PTHREAD\_PROCESS\_SHARED**

Any thread that has access to the memory where the condition variable is allocated can access the condition variable.

For more information about these attributes, see *pthread\_condattr\_getpshared()* and *pthread\_condattr\_setpshared()*.

**Returns:**

|        |                                                                                 |
|--------|---------------------------------------------------------------------------------|
| EOK    | Success.                                                                        |
| EAGAIN | All kernel synchronization objects are in use.                                  |
| EBUSY  | Previously initialized condition variable, <i>cond</i> , hasn't been destroyed. |
| EFAULT | A fault occurred when the kernel tried to access <i>cond</i> or <i>attr</i> .   |
| EINVAL | The value specified by <i>cond</i> is invalid.                                  |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_condattr\_init()*, *pthread\_cond\_destroy()*

## ***pthread\_cond\_signal()***

© 2004, QNX Software Systems Ltd.

*Unblock a thread that's waiting on a condition variable*

### **Synopsis:**

```
#include <pthread.h>

int pthread_cond_signal(pthread_cond_t* cond);
```

### **Arguments:**

*cond* A pointer to the `pthread_cond_t` object for which you want to unblock the highest-priority thread.

### **Library:**

`libc`

### **Description:**

The `pthread_cond_signal()` function unblocks the highest priority thread that's waiting on the condition variable, *cond*.

If more than one thread at the highest priority is waiting, `pthread_cond_signal()` unblocks the one that has been waiting the longest.

### **Returns:**

EOK Success.

EFAULT A fault occurred trying to access the buffers provided.

EINVAL Invalid condition variable *cond*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_cond\_broadcast()*, *pthread\_cond\_wait()*,  
*SyncCondvarSignal()*.

## ***pthread\_cond\_timedwait()***

© 2004, QNX Software Systems Ltd.

*Wait on a condition variable, with a time limit*

### **Synopsis:**

```
#include <pthread.h>
#include <time.h>

int pthread_cond_timedwait(
 pthread_cond_t* cond,
 pthread_mutex_t* mutex,
 const struct timespec* abstime);
```

### **Arguments:**

*cond*            The condition variable on which to block the thread.

*mutex*           The mutex associated with the condition variable.

*abstime*        A pointer to a **timespec** structure that specifies the maximum time to block the thread, expressed as an absolute time.

### **Library:**

**libc**

### **Description:**

The *pthread\_cond\_timedwait()* function blocks the calling thread on the condition variable *cond*, and unlocks the associated mutex *mutex*. The calling thread must have locked *mutex* before waiting on the condition variable. Upon return from the function, the mutex is again locked and owned by the calling thread.

The calling thread is blocked until either another thread performs a signal or broadcast on the condition variable, the absolute time specified by *abstime* has passed, a signal is delivered to the thread, or the thread is canceled (waiting on a condition variable is a cancellation point). In all cases, the thread reacquires the mutex before being unblocked.



---

Don't use a recursive mutex with condition variables.

---

## Returns:

|           |                                                                                                                                                                                                                                                                                                         |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EOK       | Success, or the call was interrupted by a signal.                                                                                                                                                                                                                                                       |
| EAGAIN    | Insufficient system resources are available to wait on the condition.                                                                                                                                                                                                                                   |
| EFAULT    | A fault occurred trying to access the buffers provided.                                                                                                                                                                                                                                                 |
| EINVAL    | One or more of the following is true: <ul style="list-style-type: none"><li>• One or more of <i>cond</i>, <i>mutex</i> and <i>abstime</i> is invalid.</li><li>• Concurrent waits or timed waits on <i>cond</i> used different mutexes.</li><li>• The current thread doesn't own <i>mutex</i>.</li></ul> |
| ETIMEDOUT | The time specified by <i>abstime</i> has passed.                                                                                                                                                                                                                                                        |

## Examples:

Wait five seconds while trying to acquire control over a condition variable:

```
#include <errno.h>
#include <pthread.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <time.h>
#include <unistd.h>

pthread_mutex_t m = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t c = PTHREAD_COND_INITIALIZER;

void* t(void* x)
{
 int retval;
```

```
 if (retval = pthread_mutex_lock(&m)) {
 fprintf(stderr, "pthread_mutex_lock: %s\n",
 strerror(errno));

 exit(EXIT_FAILURE);
 }

 /*
 * Let pthread_cond_timedwait() break out and try
 * to acquire mutex
 */
 fprintf(stderr, "sleeping...\n");
 sleep(30);

 if (retval = pthread_mutex_unlock(&m)) {
 fprintf(stderr, "pthread_mutex_unlock: %s\n",
 strerror(errno));

 exit(EXIT_FAILURE);
 }

 return 0;
}

int main(int argc, char* argv[])
{
 struct timespec to;
 int retval;

 fprintf(stderr, "starting...\n");

 /*
 * Here's the interesting bit; we'll wait for
 * five seconds FROM NOW when we call
 * pthread_cond_timedwait().
 */
 memset(&to, 0, sizeof to);
 to.tv_sec = time(0) + 5;
 to.tv_nsec = 0;

 if (retval = pthread_mutex_lock(&m)) {
 fprintf(stderr, "pthread_mutex_lock %s\n",
 strerror(retval));

 exit(EXIT_FAILURE);
 }

 if (retval = pthread_create(0, 0, t, 0)) {
 fprintf(stderr, "pthread_create %s\n",
```

```
 strerror(retval));

 exit(EXIT_FAILURE);
 }

 if (retval = pthread_cond_timedwait(&c, &m, &to))
 {
 fprintf(stderr, "pthread_cond_timedwait %s\n",
 strerror(retval));

 exit(EXIT_FAILURE);
 }

 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1 (Threads)

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_cond\_broadcast()*, *pthread\_cond\_init()*,  
*pthread\_cond\_signal()*, *pthread\_cond\_wait()*, *SyncCondvarWait()*,  
*TimerTimeout()*, **timespec**

# ***pthread\_cond\_wait()***

© 2004, QNX Software Systems Ltd.

*Wait on condition variable*

## **Synopsis:**

```
#include <pthread.h>

int pthread_cond_wait(pthread_cond_t* cond,
 pthread_mutex_t* mutex);
```

## **Arguments:**

*cond*     A pointer to the `pthread_cond_t` object that you want the threads to block on.

*mutex*    The mutex that you want to unlock.

## **Library:**

`libc`

## **Description:**

The `pthread_cond_wait()` function blocks the calling thread on the condition variable *cond*, and unlocks the associated mutex *mutex*. The calling thread must have locked *mutex* before waiting on the condition variable. On return from the function, the mutex is again locked and owned by the calling thread.

The calling thread is blocked until either another thread performs a signal or broadcast on the condition variable, a signal is delivered to the thread, or the thread is canceled (waiting on a condition variable is a cancellation point). In all cases, the thread reacquires the mutex before being unblocked.



---

Don't use a recursive mutex with condition variables.

---

## **Returns:**

EOK        Success, or the call was interrupted by a signal.

EAGAIN     Insufficient system resources are available to wait on the condition.



- EFAULT     A fault occurred trying to access the buffers provided.
- EINVAL     One or more of the following is true:
- One or more of *cond* or *mutex* is invalid.
  - Concurrent waits on *cond* used different mutexes.
  - The current thread doesn't own *mutex*.

## Examples:

Use condition variables to synchronize producer and consumer threads:

```
#include <stdio.h>
#include <pthread.h>

pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
pthread_cond_t cond = PTHREAD_COND_INITIALIZER;
int condition = 0;
int count = 0;

int consume(void)
{
 while(1)
 {
 pthread_mutex_lock(&mutex);
 while(condition == 0)
 pthread_cond_wait(&cond, &mutex);
 printf("Consumed %d\n", count);
 condition = 0;
 pthread_cond_signal(&cond);
 pthread_mutex_unlock(&mutex);
 }

 return(0);
}

void* produce(void * arg)
{
 while(1)
 {
 pthread_mutex_lock(&mutex);
 while(condition == 1)
 pthread_cond_wait(&cond, &mutex);
 printf("Produced %d\n", count++);
 condition = 1;
 pthread_cond_signal(&cond);
 }
}
```

```
 pthread_mutex_unlock(&mutex);
 }
 return(0);
}

int main(void)
{
 pthread_create(NULL, NULL, &produce, NULL);
 return consume();
}
```

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_broadcast()*, *pthread\_cond\_init()*,  
*pthread\_cond\_signal()*, *pthread\_cond\_timedwait()*, *SyncCondvarWait()*

## ***pthread\_condattr\_destroy()***

*Destroy a condition-variable attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_destroy(
 pthread_condattr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_condattr_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_destroy()` function destroys the condition variable attribute object *attr*. Once you've destroyed the condition-variable attribute object, don't reuse it until you've reinitialized it by calling `pthread_condattr_init()`.

### **Returns:**

EOK Success.  
EINVAL Invalid condition variable attribute object *attr*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

Interrupt handler No

*continued...*

## *pthread\_condattr\_destroy()*

© 2004, QNX Software Systems Ltd.

---

### **Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

### **See also:**

*pthread\_condattr\_init()*, *pthread\_cond\_init()*

## ***pthread\_condattr\_getclock()***

*Get the clock attribute from a condition-variable attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_getclock(
 const pthread_condattr_t * attr,
 clockid_t * id);
```

### **Arguments:**

*attr* A pointer to the `pthread_condattr_t` object from which you want to get the clock.

*id* A pointer to a `clockid_t` object where the function can store the clock ID.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_getclock()` function obtains the value of the clock attribute from the attributes object referenced by *attr*.

The clock attribute is the clock ID of the clock that's used to measure the timeout service of `pthread_cond_timedwait()`. The default value of the clock attribute refers to the system clock.

### **Returns:**

EOK Success.

EINVAL Invalid value *attr*.

### **Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_init()*, *pthread\_cond\_timedwait()*,  
*pthread\_condattr\_destroy()*, *pthread\_condattr\_getpshared()*,  
*pthread\_condattr\_init()*, *pthread\_condattr\_setclock()*,  
*pthread\_condattr\_setpshared()*, *pthread\_create()*

## ***pthread\_condattr\_getpshared()***

*Get the process-shared attribute from a condition variable attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_getpshared(
 const pthread_condattr_t* attr,
 int* pshared);
```

### **Arguments:**

*attr*            A pointer to the `pthread_condattr_t` object from which you want to get the attribute.

*pshared*        A pointer to a location where the function can store the process-shared attribute. For the possible values, see `pthread_condattr_setpshared()`.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_getpshared()` function stores, in the memory pointed to by *pshared*, the process-shared attribute from a condition variable attribute object, *attr*.

### **Returns:**

EOK            Success.

EINVAL        Invalid condition variable attribute object specified by *attr*.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_condattr\_init()*, *pthread\_condattr\_setpshared()*,  
*pthread\_create()*, *pthread\_mutex\_init()*, *pthread\_cond\_init()*.



### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_init(pthread_condattr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_condattr_t` object that you want to initialize.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_init()` function initializes the attributes in the condition variable attribute object *attr* to default values. Pass *attr* to `pthread_cond_init()` to define the attributes of the condition variable.

Condition variables have at least the following attributes defined:

#### **PTHREAD\_PROCESS\_PRIVATE**

The condition variable can be accessed only by threads created within the same process as the thread that initialized the condition variable.

#### **PTHREAD\_PROCESS\_SHARED**

Any thread that has access to the memory where the condition variable is allocated can access the condition variable.

For more information about these attributes, see `pthread_condattr_getpshared()` and `pthread_condattr_setpshared()`.

**Returns:**

|        |                                                                                         |
|--------|-----------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                |
| ENOMEM | Insufficient memory to initialize the condition variable attribute object <i>attr</i> . |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_condattr\_destroy()*, *pthread\_cond\_init()*

## ***pthread\_condattr\_setclock()***

*Set the clock attribute of a condition-variable attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_setclock(
 pthread_condattr_t * attr,
 clockid_t id);
```

### **Arguments:**

- attr* A pointer to the `pthread_condattr_t` object for which you want to set the clock. You must have initialized this object by calling `pthread_condattr_init()` before calling `pthread_condattr_setclock()`.
- id* A `clockid_t` object that specifies the ID of the clock that you want to use for the condition variable.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_setclock()` function sets the clock attribute in an initialized attributes object referenced by *attr*. If `pthread_condattr_setclock()` is called with an *id* argument that refers to a CPU-time clock, the call fails.

The clock attribute is the clock ID of the clock that you want to use to measure the timeout service of `pthread_cond_timedwait()`. The default value of the clock attribute refers to the system clock.

### **Returns:**

- EOK            Success.
- EINVAL        Invalid value *attr*.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_init()*, *pthread\_cond\_timedwait()*,  
*pthread\_condattr\_destroy()*, *pthread\_condattr\_getclock()*,  
*pthread\_condattr\_getpshared()*, *pthread\_condattr\_init()*,  
*pthread\_condattr\_setpshared()*, *pthread\_create()*

## ***pthread\_condattr\_setpshared()***

*Set the process-shared attribute in a condition variable attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_condattr_setpshared(
 pthread_condattr_t* attr,
 int pshared);
```

### **Arguments:**

- attr*            A pointer to the `pthread_condattr_t` object for which you want to set the attribute.
- pshared*        The new value of the process-shared attribute; one of:
- `PTHREAD_PROCESS_SHARED` — any thread that has access to the memory where the condition variable is allocated can operate on it, even if the condition variable is allocated in memory that's shared by multiple processes.
  - `PTHREAD_PROCESS_PRIVATE` — the condition variable can be accessed only by threads created within the same process as the thread that initialized the condition variable; if threads from other processes try to access the `PTHREAD_PROCESS_PRIVATE` condition variable, the behavior is undefined.

### **Library:**

`libc`

### **Description:**

The `pthread_condattr_setpshared()` function sets the process-shared attribute in a condition variable attribute object, *attr* to the value given by *pshared*.

The default value of the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`.

**Returns:**

- EOK Success.
- EINVAL The condition variable attribute object specified by *attr*, or the new value specified in *pshared* isn't valid.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_condattr\_init()*, *pthread\_condattr\_getpshared()*,  
*pthread\_create()*, *pthread\_mutex\_init()*, *pthread\_cond\_init()*

**Synopsis:**

```
#include <pthread.h>

int pthread_create(pthread_t* thread,
 const pthread_attr_t* attr,
 void* (*start_routine)(void*),
 void* arg);
```

**Arguments:**

*thread*            NULL, or a pointer to a **pthread\_t** object where the function can store the thread ID of the new thread.

*attr*              A pointer to a **pthread\_attr\_t** structure that specifies the attributes of the new thread. Instead of manipulating the members of this structure directly, use *pthread\_attr\_init()* and the *pthread\_attr\_set\_\** functions. For the exceptions, see “QNX extensions,” below.

If *attr* is NULL, the default attributes are used (see *pthread\_attr\_init()*).



If you modify the attributes in *attr* after creating the thread, the thread’s attributes aren’t affected.

---

*start\_routine*    The routine where the thread begins, with *arg* as its only argument. If *start\_routine()* returns, there’s an implicit call to *pthread\_exit()*, using the return value of *start\_routine()* as the exit status.

The thread in which *main()* was invoked behaves differently. When it returns from *main()*, there’s an implicit call to *exit()*, using the return value of *main()* as the exit status.

*arg*                The argument to pass to *start\_routine*.

## Library:

libc

## Description:

The *pthread\_create()* function creates a new thread, with the attributes specified in the thread attribute object *attr*. The created thread inherits the signal mask of the parent thread, and its set of pending signals is empty.



- You must call *pthread\_join()* or *pthread\_detach()* for threads created with a *detachstate* attribute of PTHREAD\_CREATE\_JOINABLE (the default) before all of the resources associated with the thread can be released at thread termination.
- If you set the *stacksize* member of *attr*, the thread's actual stack size is rounded up to a multiple of the system page size (which you can get by using the `_SC_PAGESIZE` constant in a call to *sysconf()*) if the system allocates the stack (the *stackaddr* member of *attr* is set to NULL). If the stack was previously allocated by the application, its size isn't changed.

## QNX extensions

If you adhere to the POSIX standard, there are some thread attributes that you can't specify before creating the thread:

- You can't disable cancellation for a thread.
- You can't set the thread's cancellation type.
- You can't specify what happens when a signal is delivered to the thread.

There are no *pthread\_attr\_set\_\** functions for these attributes.



As an QNX extension, you can OR the following bits into the *flags* member of the `pthread_attr_t` structure before calling `pthread_create()`:

`PTHREAD_CANCEL_ENABLE`

Cancellation requests may be acted on according to the cancellation type (the default).

`PTHREAD_CANCEL_DISABLE`

Cancellation requests are held pending.

`PTHREAD_CANCEL_ASYNCHRONOUS`

If cancellation is enabled, new or pending cancellation requests may be acted on immediately.

`PTHREAD_CANCEL_DEFERRED`

If cancellation is enabled, cancellation requests are held pending until a cancellation point (the default).

`PTHREAD_MULTISIG_ALLOW`

Terminate all the threads in the process (the POSIX default).

`PTHREAD_MULTISIG_DISALLOW`

Terminate only the thread that received the signal.

After creating the thread, you can change the cancellation properties by calling `pthread_setcancelstate()` and `pthread_setcanceltype()`.

## Returns:

|                     |                                                                                      |
|---------------------|--------------------------------------------------------------------------------------|
| <code>EAGAIN</code> | Insufficient system resources to create thread.                                      |
| <code>EFAULT</code> | An error occurred trying to access the buffers or the <i>start_routine</i> provided. |
| <code>EINVAL</code> | Invalid thread attribute object <i>attr</i> .                                        |
| <code>EOK</code>    | Success.                                                                             |

## Examples:

Create a thread in a detached state:

```
#include <stdio.h>
#include <stdlib.h>
#include <pthread.h>

void* function(void* arg)
{
 printf("This is thread %d\n", pthread_self());
 return(0);
}

int main(void)
{
 pthread_attr_t attr;

 pthread_attr_init(&attr);
 pthread_attr_setdetachstate(
 &attr, PTHREAD_CREATE_DETACHED);
 pthread_create(NULL, &attr, &function, NULL);

 /* Allow threads to run for 60 seconds. */
 sleep(60);
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1 (Threads)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_attr\_init(), pthread\_exit(), pthread\_setcancelstate(),  
pthread\_setcanceltype(), sysconf(), ThreadCreate()*

## ***pthread\_detach()***

© 2004, QNX Software Systems Ltd.

*Detach a thread from a process*

### **Synopsis:**

```
#include <pthread.h>

int pthread_detach(pthread_t thread);
```

### **Arguments:**

*thread*     The ID of the thread that you want to detach, which you can get when you call *pthread\_create()* or *pthread\_self()*.

### **Library:**

libc

### **Description:**

The *pthread\_detach()* function detaches the thread with the given ID from its process. When a detached thread terminates, all system resources allocated to that thread are immediately reclaimed.

### **Returns:**

EOK        Success.

EINVAL     The thread *thread* is already detached.

ESRCH      The thread *thread* doesn't exist.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*pthread\_join()*, *ThreadDetach()*.

## ***pthread\_equal()***

© 2004, QNX Software Systems Ltd.

*Compare two thread IDs*

### **Synopsis:**

```
#include <pthread.h>

int pthread_equal(pthread_t t1,
 pthread_t t2);
```

### **Arguments:**

*t1, t2* The thread IDs that you want to compare. You can get the IDs when you call *pthread\_create()* or *pthread\_self()*.

### **Library:**

libc

### **Description:**

The *pthread\_equal()* function compares the thread IDs of *t1* and *t2*. It doesn't check to see if they're valid thread IDs.

### **Returns:**

A nonzero value  
The *t1* and *t2* thread IDs are equal.

0 The thread IDs aren't equal.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_self()*

## ***pthread\_exit()***

© 2004, QNX Software Systems Ltd.

*Terminate a thread*

### **Synopsis:**

```
#include <pthread.h>

void pthread_exit(void* value_ptr);
```

### **Arguments:**

*value\_ptr*     A pointer to a value that you want to be made available to any thread joining the thread that you're terminating.

### **Library:**

libc

### **Description:**

The *pthread\_exit()* function terminates the calling thread. If the thread is joinable, the value *value\_ptr* is made available to any thread joining the terminating thread (only one thread can get the return status). If the thread is detached, all system resources allocated to the thread are immediately reclaimed.

Before the thread is terminated, any cancellation cleanup handlers that have been pushed are popped and executed, and any thread-specific-data destructor functions are executed. Thread termination doesn't implicitly release any process resources such as mutexes or file descriptors, or perform any process-cleanup actions such as calling *atexit()* handlers.

An implicit call to *pthread\_exit()* is made when a thread other than the thread in which *main()* was first invoked returns from the start routine that was used to create it. The return value of the start routine is used as the thread's exit status.





---

Don't call *pthread\_exit()* from cancellation-cleanup handlers or thread-specific-data destructor functions.

---

For the last process thread, *pthread\_exit()* behaves as if you called *exit(0)*.

### Classification:

POSIX 1003.1 (Threads)

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*atexit()*, *exit()*, *pthread\_create()*, *pthread\_cleanup\_push()*, *pthread\_cleanup\_pop()*, *pthread\_join()*, *ThreadDestroy()*.

# ***pthread\_getconcurrency()***

© 2004, QNX Software Systems Ltd.

*Get the level of thread concurrency*

## **Synopsis:**

```
#include <pthread.h>

int pthread_getconcurrency(void);
```

## **Library:**

libc

## **Description:**

QNX doesn't support the multiplexing of user threads on top of several kernel scheduled entities. As such, the *pthread\_setconcurrency()* and *pthread\_getconcurrency()* functions are provided for source code compatibility but they have no effect when called. To maintain the function semantics, the *new\_level* parameter is saved when *pthread\_setconcurrency()* is called so that a subsequent call to *pthread\_getconcurrency()* returns the same value.

## **Returns:**

The concurrency level set by a previous call to *pthread\_setconcurrency()*, or 0 if there was no previous call.

## **Classification:**

Standard Unix

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_setconcurrency()*

## ***pthread\_getcpuclockid()***

© 2004, QNX Software Systems Ltd.

*Return the clock ID of the CPU-time clock from a specified thread*

### **Synopsis:**

```
#include <sys/types.h>
#include <time.h>
#include <pthread.h>

extern int pthread_getcpuclockid(
 pthread_t id,
 clockid_t* clock_id);
```

### **Arguments:**

*thread*      The ID of the thread that you want to get the clock ID for, which you can get when you call *pthread\_create()* or *pthread\_self()*.

*clock\_id*    A pointer to a `clockid_t` object where the function can store the clock ID.

### **Library:**

`libc`

### **Description:**

The *pthread\_getcpuclockid()* function returns the clock ID of the CPU-time clock of the thread specified by *id*, if the thread specified by *id* exists.

### **Returns:**

0      Success.

-1     An error occurred (*errno* is set).

### **Errors:**

ESRCH    The value specified by *id* doesn't refer to an existing thread.

**Classification:**

POSIX 1003.1d (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*clock\_getcpuclockid(), clock\_getres(), clock\_gettime(), ClockId(), clock\_settime(), pthread\_getcpuclockid(), timer\_create()*

## ***pthread\_getschedparam()***

© 2004, QNX Software Systems Ltd.

*Get a thread's scheduling parameters*

### **Synopsis:**

```
#include <pthread.h>

int pthread_getschedparam(
 const pthread_t thread,
 int *policy,
 struct sched_param *param);
```

### **Arguments:**

- thread*     The ID of the thread that you want to get the scheduling parameters for. You can get a thread ID by calling *pthread\_create()* or *pthread\_self()*.
- policy*     A pointer to a location where the function can store the scheduling policy; one of SCHED\_FIFO, SCHED\_RR, SCHED\_SPORADIC, or SCHED\_OTHER.
- param*     A pointer to a **sched\_param** structure where the function can store the scheduling parameters.

### **Library:**

**libc**

### **Description:**

The *pthread\_getschedparam()* function gets the scheduling policy and associated scheduling parameters of thread *thread* and places them in *policy* and *param*.

### **Returns:**

- EOK            Success.
- EFAULT        A fault occurred trying to access the buffers provided.
- ESRCH         Invalid thread ID *thread*.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_setschedparam()*, **sched\_param**

## ***pthread\_getspecific()***

© 2004, QNX Software Systems Ltd.

*Get a thread-specific data value*

### **Synopsis:**

```
#include <pthread.h>

void* pthread_getspecific(pthread_key_t key);
```

### **Arguments:**

*key* The key associated with the data that you want to get. See *pthread\_key\_create()*.

### **Library:**

libc

### **Description:**

The *pthread\_getspecific()* function returns the thread-specific data value currently bound to the thread-specific-data key *key* in the calling thread, or NULL if no value is bound or the key doesn't exist. You can call this function from a thread-specific-data destructor function.



---

You must call this function with a key that you got from *pthread\_key\_create()*. You can't use a key after destroying it with *pthread\_key\_delete()*.

---

### **Returns:**

The data value, or NULL.

### **Examples:**

See *pthread\_key\_create()*.

### **Classification:**

POSIX 1003.1 (Threads)



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_key\_create()*, *pthread\_key\_delete()*, *pthread\_setspecific()*.

# ***pthread\_join()***

© 2004, QNX Software Systems Ltd.

*Join thread*

## **Synopsis:**

```
#include <pthread.h>

int pthread_join(pthread_t thread,
 void** value_ptr);
```

## **Arguments:**

*thread*            The target thread whose termination you're waiting for.

*value\_ptr*        NULL, or a pointer to a location where the function can store the value passed to *pthread\_exit()* by the target thread.

## **Library:**

`libc`

## **Description:**

The *pthread\_join()* function blocks the calling thread until the target thread *thread* terminates, unless *thread* has already terminated. If *value\_ptr* is non-NULL and *pthread\_join()* returns successfully, then the value passed to *pthread\_exit()* by the target thread is placed in *value\_ptr*. If the target thread has been canceled then *value\_ptr* is set to `PTHREAD_CANCELED`.



---

The non-POSIX *pthread\_timedjoin()* function is similar to *pthread\_join()*, except that an error is returned if the join doesn't occur before a given time.

---

The target thread must be joinable. Multiple *pthread\_join()*, *pthread\_timedjoin()*, *ThreadJoin()*, and *ThreadJoin\_r()* calls on the same target thread aren't allowed. When *pthread\_join()* returns successfully, the target thread has been terminated.

**Returns:**

|         |                                                         |
|---------|---------------------------------------------------------|
| EOK     | Success.                                                |
| EBUSY   | The thread <i>thread</i> is already being joined.       |
| EDEADLK | The thread <i>thread</i> is the calling thread.         |
| EFAULT  | A fault occurred trying to access the buffers provided. |
| EINTR   | The function call was interrupted.                      |
| EINVAL  | The thread <i>thread</i> isn't joinable.                |
| ESRCH   | The thread <i>thread</i> doesn't exist.                 |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_detach()*, *pthread\_exit()*,  
*pthread\_timedjoin()*, *ThreadJoin()*, *ThreadJoin\_r()*

## ***pthread\_key\_create()***

© 2004, QNX Software Systems Ltd.

*Create a thread-specific data key*

### **Synopsis:**

```
#include <pthread.h>

int pthread_key_create(pthread_key_t * key,
 void (*destructor) (void *));
```

### **Arguments:**

|                   |                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------|
| <i>key</i>        | A pointer to a <code>pthread_key_t</code> object where the function can store the new key. |
| <i>destructor</i> | NULL, or a function to be called when you destroy the key.                                 |

### **Library:**

`libc`

### **Description:**

The `pthread_key_create()` function creates a thread-specific data key that's available to all threads in the process and binds an optional destructor function `destructor` to the key. If the function completes successfully the new key is returned in `key`.



---

Although the same key may be used by different threads, the values bound to the key using `pthread_setspecific()` are maintained on a per-thread basis and persist only for the lifetime of the thread.

---

When you create a key, the value NULL is bound with the key in all active threads. When you create a thread, the value NULL is bound to all defined keys in the new thread.

You can optionally associate a destructor function with each key value. At thread exit, if the key has a non-NULL destructor pointer, and the thread has a non-NULL value bound to the key, the destructor function is called with the bound value as its only argument. The

order of destructor calls is unspecified if more than one destructor exists for a thread when it exits.

If, after all destructor functions have been called for a thread, there are still non-NULL bound values, the destructor function is called repeatedly a maximum of PTHREAD\_DESTRUCTOR\_ITERATIONS times for each non-NULL bound value.

## Returns:

|        |                                                                                    |
|--------|------------------------------------------------------------------------------------|
| EOK    | Success.                                                                           |
| EAGAIN | Insufficient system resources to create key or PTHREAD_KEYS_MAX has been exceeded. |
| ENOMEM | Insufficient memory to create key.                                                 |

## Examples:

This example shows how you can use thread-specific data to provide per-thread data in a thread-safe function:

```
#include <stdio.h>
#include <stdlib.h>
#include <pthread.h>

pthread_key_t buffer_key;

void buffer_key_destruct(void *value)
{
 free(value);
 pthread_setspecific(buffer_key, NULL);
}

char *lookup(void)
{
 char *string;

 string = (char *)pthread_getspecific(buffer_key);
 if(string == NULL) {
 string = (char *) malloc(32);
 sprintf(string, "This is thread %d\n", pthread_self());
 pthread_setspecific(buffer_key, (void *)string);
 }
}
```

```
 return(string);
 }

void *function(void *arg)
{
 while(1) {
 puts(lookup());
 }

 return(0);
}

int main(void)
{
 pthread_key_create(&buffer_key,
 &buffer_key_destruct);
 pthread_create(NULL, NULL, &function, NULL);

 /* Let the threads run for 60 seconds. */
 sleep(60);

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

The *pthread\_key\_create()* function is part of the POSIX 1003.1-1996 draft standard; its specific behavior may vary from system to system.

Before each destructor is called, the thread's value for the corresponding key is set to NULL. Calling:

```
pthread_setspecific(key, NULL);
```

in a key destructor isn't required; this lets you use the same destructor for several keys.

**See also:**

*pthread\_getspecific()*, *pthread\_setspecific()*, *pthread\_key\_delete()*

## ***pthread\_key\_delete()***

© 2004, QNX Software Systems Ltd.

*Delete a thread-specific data key*

### **Synopsis:**

```
#include <pthread.h>

int pthread_key_delete(pthread_key_t key);
```

### **Arguments:**

*key*     The key, which you created by calling *pthread\_key\_create()*, that you want to delete.

### **Library:**

`libc`

### **Description:**

The *pthread\_key\_delete()* function deletes the thread-specific data key *key* that you previously created with *pthread\_key\_create()*. The destructor function bound to *key* isn't called by this function, and won't be called at thread termination. You can call this function from a thread specific data destructor function.

If you need to release any data bound to the key in any threads, do so before deleting the key.

### **Returns:**

EOK        Success.

EINVAL     Invalid thread-specific data key *key*.

### **Classification:**

POSIX 1003.1 (Threads)



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_key\_create()*

## ***pthread\_kill()***

© 2004, QNX Software Systems Ltd.

*Send a signal to a thread*

### **Synopsis:**

```
#include <signal.h>

int pthread_kill(pthread_t thread,
 int sig);
```

### **Arguments:**

*thread*     The ID of the thread that you want to send the signal to, which you can get when you call *pthread\_create()* or *pthread\_self()*.

*sig*        The signal that you want to send, or 0 if you just want to check for errors.

### **Library:**

libc

### **Description:**

The *pthread\_kill()* function sends the signal *sig* to the thread *thread*. The target thread and calling thread must be in the same process. If *sig* is zero, error checking is performed but no signal is sent.

### **Returns:**

EOK         Success.

EAGAIN      Insufficient system resources are available to deliver the signal.

ESRCH       Invalid thread ID *thread*.

EINVAL      Invalid signal number *sig*.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*kill()*, *ThreadDestroy()*

## ***pthread\_mutex\_destroy()***

© 2004, QNX Software Systems Ltd.

*Destroy a mutex*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutex_destroy(pthread_mutex_t* mutex);
```

### **Arguments:**

*mutex*     A pointer to the `pthread_mutex_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_mutex_destroy()` function destroys the unlocked mutex *mutex*.

You can only destroy a locked mutex provided you're the owner of that mutex.



---

Once you've destroyed a mutex, don't reuse it without reinitializing it by calling `pthread_mutex_init()`.

---

### **Returns:**

EOK        Success.

EBUSY     The *mutex* is locked by another thread.

EINVAL    Invalid mutex *mutex*.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_init()*, *SyncDestroy()*, *SyncMutexRevive()*

## ***pthread\_mutex\_getprioceiling()***

© 2004, QNX Software Systems Ltd.

*Get a mutex's priority ceiling*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutex_getprioceiling(
 const pthread_mutex_t* mutex,
 int* prioceiling);
```

### **Arguments:**

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <i>mutex</i>       | A pointer to the <code>pthread_mutex_t</code> object that you want to priority ceiling for. |
| <i>prioceiling</i> | A pointer to a location where the function can store the priority ceiling.                  |

### **Library:**

`libc`

### **Description:**

The `pthread_mutex_getprioceiling()` function returns the priority ceiling of *mutex* in *prioceiling*.

### **Returns:**

|        |                                                                         |
|--------|-------------------------------------------------------------------------|
| EOK    | Success.                                                                |
| EINVAL | The mutex specified by <i>mutex</i> doesn't currently exist.            |
| EPERM  | The calling thread doesn't have permission to get the priority ceiling. |

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_setprioceiling()*

## ***pthread\_mutex\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize mutex*

### **Synopsis:**

```
#include <pthread.h>

pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;

int pthread_mutex_init(
 pthread_mutex_t* mutex,
 const pthread_mutexattr_t* attr);
```

### **Arguments:**

*mutex* A pointer to the `pthread_mutex_t` object that you want to initialize.

*attr* NULL, or a pointer to a `pthread_mutexattr_t` object that specifies the attributes that you want to use for the mutex. For more information, see *pthread\_mutexattr\_init()*.

### **Library:**

`libc`

### **Description:**

The *pthread\_mutex\_init()* function initializes the given mutex object, using the attributes specified by the mutex attributes object *attr*. If *attr* is NULL then the mutex is initialized with the default attributes (see *pthread\_mutexattr\_init()*). After initialization, the mutex is in an unlocked state.

You can initialize a statically allocated mutex with the default attributes by assigning to it the macro `PTHREAD_MUTEX_INITIALIZER` or `PTHREAD_RMUTEX_INITIALIZER` (for recursive mutexes).



## Returns:

|        |                                                                                |
|--------|--------------------------------------------------------------------------------|
| EOK    | Success.                                                                       |
| EAGAIN | All kernel synchronization objects are in use.                                 |
| EBUSY  | The given mutex was previously initialized and hasn't been destroyed.          |
| EFAULT | A fault occurred when the kernel tried to access <i>mutex</i> or <i>attr</i> . |
| EINVAL | The value specified by <i>attr</i> is invalid.                                 |

## Classification:

POSIX 1003.1 (Threads)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_mutexattr\_init()*, *pthread\_mutex\_destroy()*, *SyncTypeCreate()*

## ***pthread\_mutex\_lock()***

© 2004, QNX Software Systems Ltd.

*Lock a mutex*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutex_lock(pthread_mutex_t* mutex);
```

### **Arguments:**

*mutex*     A pointer to the `pthread_mutex_t` object that you want to lock.

### **Library:**

`libc`

### **Description:**

The `pthread_mutex_lock()` function locks the mutex object referenced by *mutex*. If the mutex is already locked, then the calling thread blocks until it has acquired the mutex. When the function returns, the mutex object is locked and owned by the calling thread.

If the mutex allows recursive behavior, a call to `pthread_mutex_lock()` while you own the mutex succeeds. You can allow recursive behavior by:

- statically initializing the mutex to `PTHREAD_MUTEX_INITIALIZER`

Or:

- using `pthread_mutexattr_setrecursive()` to set the attribute to `PTHREAD_MUTEX_RECURSIVE_ALLOW` before calling `pthread_mutex_init()`.

If the mutex is recursive, you must call `pthread_mutex_unlock()` for each corresponding call to lock the mutex. The default POSIX behavior doesn't allow recursive mutexes, and returns `EDEADLK`.

If a signal is delivered to a thread that's waiting for a mutex, the thread resumes waiting for the mutex on returning from the signal handler.

**Returns:**

|         |                                                                                                |
|---------|------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                       |
| EAGAIN  | Insufficient system resources available to lock the mutex.                                     |
| EDEADLK | The calling thread already owns <i>mutex</i> , and the mutex doesn't allow recursive behavior. |
| EINVAL  | Invalid mutex <i>mutex</i> .                                                                   |

**Examples:**

This example shows how you can use a mutex to synchronize access to a shared variable. In this example, *function1()* and *function2()* both attempt to access and modify the global variable *count*. Either thread could be interrupted between modifying *count* and assigning its value to the local *tmp* variable. Locking *mutex* prevents this from happening; if one thread has *mutex* locked, the other thread waits until it's unlocked, before continuing.

```
#include <stdio.h>
#include <pthread.h>

pthread_mutex_t mutex = PTHREAD_MUTEX_INITIALIZER;
int count = 0;

void* function1(void* arg)
{
 int tmp = 0;

 while(1) {
 pthread_mutex_lock(&mutex);
 tmp = count++;
 pthread_mutex_unlock(&mutex);
 printf("Count is %d\n", tmp);

 /* snooze for 1 second */
 sleep(1);
 }

 return 0;
}
```

```
void* function2(void* arg)
{
 int tmp = 0;

 while(1) {
 pthread_mutex_lock(&mutex);
 tmp = count--;
 pthread_mutex_unlock(&mutex);
 printf("*** Count is %d\n", tmp);

 /* snooze for 2 seconds */
 sleep(2);
 }

 return 0;
}

int main(void)
{
 pthread_create(NULL, NULL, &function1, NULL);
 pthread_create(NULL, NULL, &function2, NULL);

 /* Let the threads run for 60 seconds. */
 sleep(60);

 return 0;
}
```

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_init()*, *pthread\_mutexattr\_setrecursive()*,  
*pthread\_mutex\_trylock()*, *pthread\_mutex\_unlock()*, *SyncMutexLock()*

## ***pthread\_mutex\_setprioceiling()***

© 2004, QNX Software Systems Ltd.

*Set a mutex's priority ceiling*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutex_setprioceiling(
 pthread_mutex_t* mutex,
 int prioceiling,
 int* old_ceiling);
```

### **Arguments:**

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <i>mutex</i>       | A pointer to the <code>pthread_mutex_t</code> object that you want to priority ceiling for. |
| <i>prioceiling</i> | The new value for the priority ceiling.                                                     |
| <i>old_ceiling</i> | A pointer to a location where the function can store the old value.                         |

### **Library:**

`libc`

### **Description:**

The `pthread_mutex_setprioceiling()` function locks *mutex* (or blocks until it can lock it), changes its priority ceiling to *prioceiling*, and releases it. When the change is successful, the previous priority ceiling is returned in *old\_ceiling*.

### **Returns:**

|        |                                                                                                                               |
|--------|-------------------------------------------------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                                                      |
| EINVAL | The mutex specified by <i>mutex</i> doesn't currently exist, or the priority requested by <i>prioceiling</i> is out of range. |
| EPERM  | The calling thread doesn't have permission to set the priority ceiling.                                                       |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_getprioceiling()*

# ***pthread\_mutex\_timedlock()***

© 2004, QNX Software Systems Ltd.

*Attempt to lock mutex*

## **Synopsis:**

```
#include <pthread.h>
#include <time.h>

int pthread_mutex_timedlock(
 pthread_mutex_t * mutex,
 const struct timespec * abs_timeout);
```

## **Arguments:**

|                    |                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <i>mutex</i>       | The mutex that you want to lock.                                                                                                   |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> structure that specifies the maximum time to wait to lock the mutex, expressed as an absolute time. |

## **Library:**

**libc**

## **Description:**

The *pthread\_mutex\_timedlock()* function is called to lock the mutex object referenced by *mutex*. If the mutex is already locked, the calling thread blocks until the mutex becomes available as in the *pthread\_mutex\_lock* function. If the mutex can't be locked without waiting for another thread to unlock the mutex, the wait is terminated when the specified timeout expires.

The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the clock on which timeouts are based (i.e., when the value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

The timeout is based on the **CLOCK\_REALTIME** clock. The **timespec** datatype is defined in the **<time.h>** header.



If the mutex can be locked immediately, the validity of the *abs\_timeout* parameter isn't checked, and the function won't fail with a timeout.

As a consequence of the priority inheritance rules (for mutexes initialized with the PRIO\_INHERIT protocol), if a timed mutex wait is terminated because its timeout expires, the priority of the owner of the mutex is adjusted as necessary to reflect the fact that this thread is no longer among the threads waiting for the mutex.

### Returns:

Zero on success, or an error number to indicate the error.

### Errors:

|           |                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN    | The mutex couldn't be acquired because the maximum number of recursive locks for the mutex has been exceeded.                                                                                                                                                                                                                                                                                                                                  |
| EDEADLK   | The current thread already owns the mutex.                                                                                                                                                                                                                                                                                                                                                                                                     |
| EINVAL    | The mutex was created with the protocol attribute having the value PTHREAD_PRIO_PROTECT and the calling thread's priority is higher than the mutex' current priority ceiling; the process or thread would have blocked, and the <i>abs_timeout</i> parameter specified a nanoseconds field value less than zero or greater than or equal to 1000 million; or the value specified by <i>mutex</i> doesn't refer to an initialized mutex object. |
| ETIMEDOUT | The mutex couldn't be locked before the specified timeout expired.                                                                                                                                                                                                                                                                                                                                                                             |

### Classification:

POSIX 1003.1d (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_destroy()*, *pthread\_mutex\_lock()*,  
*pthread\_mutex\_trylock()*, *pthread\_mutex\_unlock()*, **timespec**

**Synopsis:**

```
#include <pthread.h>

int pthread_mutex_trylock(pthread_mutex_t* mutex);
```

**Arguments:**

*mutex* A pointer to the `pthread_mutex_t` object that you want to try to lock.

**Library:**

`libc`

**Description:**

The `pthread_mutex_trylock()` function attempts to lock the mutex *mutex*, but doesn't block the calling thread if the mutex is already locked.

**Returns:**

|        |                                                     |
|--------|-----------------------------------------------------|
| EOK    | Success.                                            |
| EAGAIN | Insufficient resources available to lock the mutex. |
| EBUSY  | The <i>mutex</i> was already locked.                |
| EINVAL | Invalid mutex <i>mutex</i> .                        |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

Cancellation point No  
*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_mutex\_lock()*, *pthread\_mutex\_unlock()*

**Synopsis:**

```
#include <pthread.h>

int pthread_mutex_unlock(pthread_mutex_t* mutex);
```

**Arguments:**

*mutex*      A pointer to the `pthread_mutex_t` object that you want to unlock.

**Library:**

`libc`

**Description:**

The `pthread_mutex_unlock()` function unlocks the mutex *mutex*. The mutex should be owned by the calling thread. If there are threads blocked on the mutex then the highest priority waiting thread is unblocked and becomes the next owner of the mutex.

If *mutex* has been locked more than once, it must be unlocked the same number of times before the next thread is given ownership of the mutex. This only works for recursive mutexes.

**Returns:**

|        |                                           |
|--------|-------------------------------------------|
| EOK    | Success.                                  |
| EINVAL | Invalid mutex <i>mutex</i> .              |
| EPERM  | Current thread doesn't own <i>mutex</i> . |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_lock()*, *pthread\_mutex\_trylock()*, *SyncMutexUnlock()*

## ***pthread\_mutexattr\_destroy()***

*Destroy mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_destroy(
 pthread_mutexattr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_mutexattr_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_destroy()` function destroys the mutex attribute object *attr*.



---

Once you've destroyed a mutex attribute object, don't reuse it without reinitializing it by calling `pthread_mutexattr_init()`.

---

### **Returns:**

EOK Success.  
EINVAL Invalid mutex attribute object *attr*.

### **Classification:**

POSIX 1003.1 (Threads)

#### **Safety**

---

Cancellation point No

*continued...*

## *pthread\_mutexattr\_destroy()*

© 2004, QNX Software Systems Ltd.

---

### **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

### **See also:**

*pthread\_mutexattr\_init()*, *pthread\_mutex\_init()*



## ***pthread\_mutexattr\_getprioceiling()***

*Get the priority ceiling of a mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_getprioceiling(
 const pthread_mutexattr_t* attr,
 int* prioceiling);
```

### **Arguments:**

*attr*            A pointer to the `pthread_mutexattr_t` object that you want to get the attribute from.

*prioceiling*    A pointer to a location where the function can store the priority ceiling.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_getprioceiling()` function sets *prioceiling* to the current mutex attribute *attr*'s scheduling priority ceiling. The mutex attribute object *attr* must have been previously created with `pthread_mutexattr_init()`.

### **Returns:**

EOK            Success.

EINVAL        Invalid value specified by *attr* or *prioceiling*.

EPERM        The caller doesn't have the privilege to perform the operation.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_init()*, *pthread\_create()*, *pthread\_mutex\_init()*,  
*pthread\_mutexattr\_getprotocol()*, *pthread\_mutexattr\_getrecursive()*,  
*pthread\_mutexattr\_setprioceiling()*, *pthread\_mutexattr\_setrecursive()*

**Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_getprotocol(
 pthread_mutexattr * attr,
 int * protocol);
```

**Arguments:**

*attr* A pointer to the `pthread_mutexattr_t` object that you want to get the attribute from.

*protocol* A pointer to a location where the function can store the scheduling protocol.

**Library:**

`libc`

**Description:**

The `pthread_mutexattr_getprotocol()` function sets *protocol* to the current mutex attribute *attr*'s scheduling protocol. The structure pointed to by *attr* must have been previously created with `pthread_mutexattr_init()`.

The *protocol* attribute defines the protocol for using mutexes: Currently, *protocol* may be set to:

**PTHREAD\_PRIO\_INHERIT**

When a thread is blocking higher-priority threads by locking one or more mutexes with this attribute, the thread's priority is raised to that of the highest priority thread waiting on the PTHREAD\_PRIO\_INHERIT mutex.

**PTHREAD\_PRIO\_PROTECT**

The thread executes at the highest priority or priority ceilings of all the mutexes owned by the thread and initialized with

PTHREAD\_PRIO\_PROTECT, whether other threads are blocked or not.

A thread holding a PTHREAD\_PRIO\_INHERIT mutex won't be moved to the tail of the scheduling queue if its original priority is changed (by a call to *pthread\_schedsetparam()*, for example). This remains true if the thread unlocks the PTHREAD\_PRIO\_INHERIT mutex.



---

The POSIX *protocol* of PTHREAD\_PRIO\_NONE isn't currently supported.

---

### Returns:

EOK            Success.  
EINVAL        Invalid mutex attribute *attr*.

### Classification:

POSIX 1003.1 (Threads)

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_mutexattr\_setprotocol()*, *pthread\_mutexattr\_setrecursive()*

## ***pthread\_mutexattr\_getpshared()***

*Get the process-shared attribute from a mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_getpshared(
 const pthread_mutexattr_t* attr,
 int* pshared);
```

### **Arguments:**

*attr*            A pointer to the `pthread_mutexattr_t` object that you want to get the attribute from.

*pshared*        A pointer to a location where the function can store the process-shared attribute.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_getpshared()` function gets the process-shared attribute from the mutex attribute object *attr* and stores it in the memory pointed to by *pshared*.

If the process-shared attribute is set to `PTHREAD_PROCESS_SHARED`, any thread that has access to the memory where the condition variable is allocated can operate on it, even if the condition variable is allocated in memory that's shared by multiple processes.

If the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`, the condition variable can only be accessed by threads created within the same process as the thread that initialized the condition variable; if threads from other processes try to access the `PTHREAD_PROCESS_PRIVATE` condition variable, the behavior is undefined. The default value of the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`.

**Returns:**

- EOK        Success.
- EINVAL     The mutex attribute object specified by *attr* is invalid.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*pthread\_cond\_init()*, *pthread\_create()*, *pthread\_mutex\_init()*,  
*pthread\_mutexattr\_setpshared()*, *pthread\_mutexattr\_setrecursive()*

## ***pthread\_mutexattr\_getrecursive()***

*Get the recursive attribute from a mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_getrecursive(
 const pthread_mutexattr_t* attr,
 int* recursive);
```

### **Arguments:**

*attr*            A pointer to the `pthread_mutexattr_t` object that you want to get the attribute from.

*pshared*        A pointer to a location where the function can store the recursive attribute.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_getrecursive()` function gets the recursive attribute from the mutex attribute object *attr* and stores it in *recursive*.

If the recursive attribute is set to `PTHREAD_RECURSIVE_ENABLE`, a thread that has already locked the mutex can lock it again without blocking. If the recursive attribute is set to `PTHREAD_RECURSIVE_DISABLE`, any thread that tries to lock the mutex will block, if that mutex is already locked.

The default value of the recursive attribute is `PTHREAD_RECURSIVE_DISABLE`.

### **Returns:**

`EOK`            Success.

`EINVAL`        Invalid mutex attribute object *attr*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutexattr\_init()*, *pthread\_mutexattr\_setrecursive()*



### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_gettype(
 const pthread_mutexattr_t * attr,
 int * type);
```

### **Arguments:**

*attr* A pointer to the `pthread_mutexattr_t` object that you want to get the attribute from.

*type* A pointer to a location where the function can store the type.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_gettype()` function gets the mutex type attribute in the *type* parameter. Valid mutex types include:

#### **PTHREAD\_MUTEX\_NORMAL**

No deadlock detection. A thread that attempts to relock this mutex without first unlocking it deadlocks. Attempts to unlock a mutex locked by a different thread or attempts to unlock an unlocked mutex result in undefined behavior.

#### **PTHREAD\_MUTEX\_ERRORCHECK**

Provides error checking. A thread returns with an error when it attempts to:

- Relock this mutex without first unlocking it.
- Unlock a mutex that another thread has locked.
- Unlock an unlocked mutex.

PTHREAD\_MUTEX\_RECURSIVE

A thread that attempts to relock this mutex without first unlocking it succeeds in locking the mutex. The relocking deadlock that can occur with mutexes of type PTHREAD\_MUTEX\_NORMAL can't occur with this mutex type. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread that attempts to unlock a mutex that another thread has locked, or unlock an unlocked mutex, returns with an error.

PTHREAD\_MUTEX\_DEFAULT

The default value of the *type* attribute. Attempts to recursively lock a mutex of this type, or unlock a mutex of this type that isn't locked by the calling thread, or unlock a mutex of this type that isn't locked, results in undefined behavior.

**Returns:**

Zero, and the value of the *type* attribute of *attr* is stored in the object referenced by the *type* parameter; otherwise, an error.

**Errors:**

EINVAL Invalid value specified by *attr*.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **Caveats:**

An application shouldn't use a PTHREAD\_MUTEX\_RECURSIVE mutex with condition variables because the implicit unlock performed for a *pthread\_cond\_wait()* or *pthread\_cond\_timedwait()* may not actually release the mutex (if it's been locked multiple times). If this happens, no other thread can satisfy the condition of the predicate.

### **See also:**

*pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*,  
*pthread\_mutexattr\_settype()*

## ***pthread\_mutexattr\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize a mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_init(
 const pthread_mutexattr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_mutexattr_t` object that you want to initialize.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_init()` function initializes the attributes in the mutex attribute object *attr* to their default values. After initializing a mutex attribute object, you can use it to initialize one or more mutexes by calling `pthread_mutex_init()`.

The mutex attributes and their default values are:

*protocol*      PTHREAD\_PRIO\_INHERIT  
*recursive*     PTHREAD\_RECURSIVE\_DISABLE

After calling this function, you can use the `pthread_mutexattr_*` family of functions to make any changes to the attributes:

| <b>Get</b>                                        | <b>Set</b>                                        |
|---------------------------------------------------|---------------------------------------------------|
| <code>pthread_mutexattr_getprioceiling()</code> , | <code>pthread_mutexattr_setprioceiling()</code> , |
| <code>pthread_mutexattr_getprotocol()</code> ,    | <code>pthread_mutexattr_setprotocol()</code> ,    |

*continued...*

| <b>Get</b>                                | <b>Set</b>                                |
|-------------------------------------------|-------------------------------------------|
| <i>pthread_mutexattr_getpshared()</i> ,   | <i>pthread_mutexattr_setpshared()</i> ,   |
| <i>pthread_mutexattr_getrecursive()</i> , | <i>pthread_mutexattr_setrecursive()</i> , |
| <i>pthread_mutexattr_gettype()</i> ,      | <i>pthread_mutexattr_settype()</i>        |

**Returns:**

|        |                                                                        |
|--------|------------------------------------------------------------------------|
| EOK    | Success.                                                               |
| ENOMEM | Insufficient memory to initialize mutex attribute object <i>attr</i> . |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_init()*, *pthread\_mutexattr\_destroy()*,  
*pthread\_mutexattr\_getprioceiling()*, *pthread\_mutexattr\_getprotocol()*,  
*pthread\_mutexattr\_getpshared()*, *pthread\_mutexattr\_getrecursive()*,  
*pthread\_mutexattr\_gettype()*, *pthread\_mutexattr\_setprioceiling()*,  
*pthread\_mutexattr\_setprotocol()*, *pthread\_mutexattr\_setpshared()*,  
*pthread\_mutexattr\_setrecursive()*, *pthread\_mutexattr\_settype()*

## ***pthread\_mutexattr\_setprioceiling()*** © 2004, QNX Software Systems Ltd.

*Set the priority ceiling of a mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_setprioceiling(
 pthread_mutexattr_t* attr,
 int prioceiling);
```

### **Arguments:**

*attr*            A pointer to the `pthread_mutexattr_t` object that you want to set the attribute in.

*prioceiling*    The new value for the priority ceiling.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_setprioceiling()` function sets the mutex attribute *attr*'s scheduling priority ceiling to *prioceiling*. Note that *attr* must have been previously created with `pthread_mutexattr_init()`.

### **Returns:**

EOK            Success.

EINVAL        Invalid value specified by *attr* or *prioceiling*.

EPERM         The caller doesn't have the privilege to perform the operation.

### **Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_mutex\_init()*, *pthread\_cond\_init()*,  
*pthread\_mutexattr\_getprioceiling()*, *pthread\_mutexattr\_getprotocol()*,  
*pthread\_mutexattr\_getrecursive()*

## ***pthread\_mutexattr\_setprotocol()***

© 2004, QNX Software Systems Ltd.

*Set a mutex's scheduling protocol*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_setprotocol(
 pthread_mutexattr * attr,
 int protocol);
```

### **Arguments:**

- attr*            A pointer to the `pthread_mutexattr_t` object that you want to set the attribute in.
- protocol*        The new value of the scheduling protocol; one of:
- `PTHREAD_PRIO_INHERIT` — when a thread is blocking higher-priority threads by locking one or more mutexes with this attribute, raise the thread's priority to that of the highest priority thread waiting on the `PTHREAD_PRIO_INHERIT` mutex.
  - `PTHREAD_PRIO_PROTECT` — execute the thread at the highest priority or priority ceilings of all the mutexes owned by the thread and initialized with `PTHREAD_PRIO_PROTECT`, whether other threads are blocked or not.



---

The POSIX protocol of `PTHREAD_PRIO_NONE` isn't currently supported.

---

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_setprotocol()` function sets the mutex attribute *attr*'s scheduling protocol to *protocol*. The structure pointed to by *attr* must have been previously created with `pthread_mutexattr_init()`.



The *protocol* attribute defines the protocol for using mutexes. A thread holding a PTHREAD\_PRIO\_INHERIT mutex won't be moved to the tail of the scheduling queue if its original priority is changed (by a call to *pthread\_schedsetparam()*, for example). This remains true if the thread unlocks the PTHREAD\_PRIO\_INHERIT mutex.

**Returns:**

|         |                                                                     |
|---------|---------------------------------------------------------------------|
| EOK     | Success.                                                            |
| ENOTSUP | The <i>protocol</i> argument is an unsupported or an invalid value. |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***pthread\_mutexattr\_getprotocol()*, *pthread\_mutexattr\_getrecursive()*

## ***pthread\_mutexattr\_setpshared()***

© 2004, QNX Software Systems Ltd.

*Set the process-shared attribute in mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_setpshared(
 pthread_mutexattr_t* attr,
 int pshared);
```

### **Arguments:**

- attr*            A pointer to the `pthread_mutexattr_t` object that you want to set the attribute in.
- pshared*        The new value of the process-shared attribute; one of:
- `PTHREAD_PROCESS_SHARED` — any thread that has access to the memory where the mutex is allocated can operate on it, even if the mutex is allocated in memory that's shared by multiple processes.
  - `PTHREAD_PROCESS_PRIVATE` — the mutex can be accessed only by threads created within the same process as the thread that initialized the mutex; if threads from other processes try to access the `PTHREAD_PROCESS_PRIVATE` mutex, the behavior is undefined.

The default value of the process-shared attribute is `PTHREAD_PROCESS_PRIVATE`.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_setpshared()` function sets the process-shared attribute in a mutex attribute object, *attr*, to the value given by *pshared*.

**Returns:**

- EOK           Success.
- EINVAL        Invalid mutex attribute object, *attr*.
- EINVAL        The new value specified in *pshared* isn't  
PTHREAD\_PROCESS\_SHARED or  
PTHREAD\_PROCESS\_PRIVATE.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

- Cancellation point   No
- Interrupt handler    No
- Signal handler       Yes
- Thread               Yes

**See also:**

*pthread\_cond\_init()*, *pthread\_create()*,  
*pthread\_mutexattr\_getpshared()*, *pthread\_mutexattr\_getrecursive()*,  
*pthread\_mutex\_init()*, *pthread\_mutexattr\_setrecursive()*

## ***pthread\_mutexattr\_setrecursive()*** © 2004, QNX Software Systems Ltd.

*Set the recursive attribute in mutex attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_setrecursive(
 pthread_mutexattr_t* attr,
 int recursive);
```

### **Arguments:**

- attr*            A pointer to the `pthread_mutexattr_t` object that you want to set the attribute in.
- recursive*       The new value for the recursive attribute; one of:
- `PTHREAD_RECURSIVE_ENABLE` — a thread that has already locked the mutex can lock it again without blocking.
  - `PTHREAD_RECURSIVE_DISABLE` — any thread that tries to lock the mutex blocks, if that mutex is already locked.

The default value of the recursive attribute is `PTHREAD_RECURSIVE_DISABLE`.

### **Library:**

`libc`

### **Description:**

The `pthread_mutexattr_setrecursive()` function sets the recursive attribute in a mutex attribute object, *attr*.

### **Returns:**

- `EOK`            Success.
- `EINVAL`        Invalid mutex attribute object, *attr*, or the value specified by *recursive* isn't `PTHREAD_RECURSIVE_ENABLE` or `PTHREAD_RECURSIVE_DISABLE`.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_mutexattr\_getrecursive()*, *pthread\_mutexattr\_init()*

## ***pthread\_mutexattr\_settype()***

© 2004, QNX Software Systems Ltd.

*Set a mutex type*

### **Synopsis:**

```
#include <pthread.h>

int pthread_mutexattr_settype(
 pthread_mutexattr_t * attr,
 int type);
```

### **Arguments:**

- attr*     A pointer to the `pthread_mutexattr_t` object that you want to set the attribute in.
- type*     The new type; one of:
- `PTHREAD_MUTEX_NORMAL` — no deadlock detection. A thread that attempts to relock this mutex without first unlocking it deadlocks. Attempts to unlock a mutex locked by a different thread or attempts to unlock an unlocked mutex result in undefined behavior.
  - `PTHREAD_MUTEX_ERRORCHECK` — provides error checking. A thread returns with an error when it attempts to relock this mutex without first unlocking it, unlock a mutex that another thread has locked, or unlock an unlocked mutex.
  - `PTHREAD_MUTEX_RECURSIVE` — a thread that attempts to relock this mutex without first unlocking it succeeds in locking the mutex. The relocking deadlock that can occur with mutexes of type `PTHREAD_MUTEX_NORMAL` can't occur with this mutex type. Multiple locks of this mutex require the same number of unlocks to release the mutex before another thread can acquire the mutex. A thread that attempts to unlock a mutex that another thread has locked, or unlock an unlocked mutex, returns with an error.
  - `PTHREAD_MUTEX_DEFAULT` — the default value of the *type* attribute. Attempts to recursively lock a mutex of this type, or unlock a mutex of this type that isn't locked by

the calling thread, or unlock a mutex of this type that isn't locked, results in undefined behavior.

**Library:**

`libc`

**Description:**

The *pthread\_mutexattr\_settype()* function sets the mutex type attribute in the *type* parameter.

**Returns:**

Zero for success, or an error number.

**Errors:**

EINVAL     The value specified by *attr* or *type* is invalid.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

An application shouldn't use a PTHREAD\_MUTEX\_RECURSIVE mutex with condition variables because the implicit unlock performed for a *pthread\_cond\_wait()* or *pthread\_cond\_timedwait()* may not

actually release the mutex (if it's been locked multiple times). If this happens, no other thread can satisfy the condition of the predicate.

**See also:**

*pthread\_cond\_timedwait()*, *pthread\_cond\_wait()*,  
*pthread\_mutexattr\_gettype()*



## Synopsis:

```
#include <pthread.h>

pthread_once_t once_control = PTHREAD_ONCE_INIT;

int pthread_once(pthread_once_t* once_control,
 void (*init_routine)(void));
```

## Arguments:

*once\_control*     A pointer to a **pthread\_once\_t** object that the function uses to determine whether or not to run the initialization code.



---

You must set the **pthread\_once\_t** object to the macro **PTHREAD\_ONCE\_INIT** before using it for the first time.

---

*init\_routine*     The function that you want to call to do any required initialization.

## Library:

**libc**

## Description:

The *pthread\_once()* function uses the once-control object *once\_control* to determine whether the initialization routine *init\_routine* should be called. The first call to *pthread\_once()* by any thread in a process, with a given *once\_control*, calls *init\_routine* with no arguments. Subsequent calls of *pthread\_once()* with the same *once\_control* won't call *init\_routine*.



---

No thread will execute past this function until the *init\_routine* returns.

---

**Returns:**

- EOK        Success.
- EINVAL     Uninitialized once-control object *once\_control*.

**Examples:**

This example shows how you can use once-initialization to initialize a library; both *library\_entry\_point1()* and *library\_entry\_point2()* need to initialize the library, but that needs to happen only once:

```
#include <stdio.h>
#include <pthread.h>

pthread_once_t once_control = PTHREAD_ONCE_INIT;

void library_init(void)
{
 /* initialize the library */
}

void library_entry_point1(void)
{
 pthread_once(&once_control, library_init);

 /* do stuff for library_entry_point1... */
}

void library_entry_point2(void)
{
 pthread_once(&once_control, library_init);

 /* do stuff for library_entry_point1... */
}
```

This initializes the library once; if multiple threads call *pthread\_once()*, only one actually enters the *library\_init()* function. The other threads block at the *pthread\_once()* call until *library\_init()* has returned. The *pthread\_once()* function also ensures that *library\_init()* is only ever called once; subsequent calls to the library entry points skip the call to *library\_init()*.

## Classification:

POSIX 1003.1 (Threads)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## ***pthread\_rwlock\_destroy()***

© 2004, QNX Software Systems Ltd.

*Destroy a read-write lock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_destroy(pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_destroy()` function destroys the read-write lock referenced by *rwl*, and releases the system resources used by the lock. You can destroy the read-write lock if one of the following is true:

- no thread has a active shared or exclusive lock on *rwl*
- the calling thread has an active exclusive lock on *rwl*.



---

After successfully destroying a read-write lock, don't use it again without reinitializing it by calling `pthread_rwlock_init()`.

---

### **Returns:**

EOK Success.

EBUSY The read-write lock *rwl* is still in use. The calling thread doesn't have an exclusive lock.

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_rwlock\_init()*, *pthread\_rwlock\_rdlock()*,  
*pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
*pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*

## ***pthread\_rwlock\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize a read-write lock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_init(
 pthread_rwlock_t * rwl,
 const pthread_rwlockattr_t * attr);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to initialize.

*attr* NULL, or a pointer to a `pthread_rwlockattr_t` object that specifies the attributes you want to use for the read-write lock; see *pthread\_rwlockattr\_init()*.

### **Library:**

`libc`

### **Description:**

The *pthread\_rwlock\_init()* function initializes the read-write lock referenced by *rwl* with the attributes of *attr*. You must initialize read-write locks before using them. If *attr* is NULL, *rwl* is initialized with the default values for the attributes.

Following a successful call to *pthread\_rwlock\_init()*, the read-write lock is unlocked, and you can use it in subsequent calls to *pthread\_rwlock\_destroy()*, *pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*, and *pthread\_rwlock\_wrlock()*. This lock remains usable until you destroy it by calling *pthread\_rwlock\_destroy()*.

If the read-write lock is statically allocated, you can initialize it with the default values by setting it to `PTHREAD_RWLOCK_INITIALIZER`.

More than one thread may hold a shared lock at any time, but only one thread may hold an exclusive lock. This avoids reader and writer starvation during frequent contention by:

- favoring blocked readers over writers after a writer has just released an exclusive lock, and
- favoring writers over readers when there are no blocked readers.

Under heavy contention, the lock alternates between a single exclusive lock followed by a batch of shared locks.

### Returns:

|        |                                                                                  |
|--------|----------------------------------------------------------------------------------|
| EOK    | Success.                                                                         |
| EAGAIN | Insufficient system resources to initialize the read-write lock.                 |
| EBUSY  | The read-write lock <i>rwl</i> has been initialized or unsuccessfully destroyed. |
| EFAULT | A fault occurred when the kernel tried to access <i>rwl</i> or <i>attr</i> .     |
| EINVAL | Invalid read-write lock attribute object <i>attr</i> .                           |

### Classification:

Standard Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Beware of *priority inversion* when using read-write locks. A high-priority thread may be blocked waiting on a read-write lock locked by a low-priority thread.

The microkernel has no knowledge of read-write locks, and therefore can't boost the low-priority thread to prevent the priority inversion.

**See also:**

*pthread\_rwlockattr\_init()*, *pthread\_rwlock\_destroy()*,  
*pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_tryrdlock()*,  
*pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_wrlock()*,  
*pthread\_rwlock\_unlock()*



### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_rdlock(pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to lock for reading.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_rdlock()` function acquires a shared lock on the read-write lock referenced by *rwl*. If the read-write lock is already exclusively locked, the calling thread blocks until the exclusive lock is released.

If a signal is delivered to a thread waiting to lock a read-write lock, it will resume waiting for the lock after returning from the signal handler.

A thread may hold several read locks on the same read-write lock; it must call `pthread_rwlock_unlock()` multiple times to release its read lock.

### **Returns:**

|         |                                                                                                                                      |
|---------|--------------------------------------------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                                                             |
| EAGAIN  | On the first use of statically initialized read-write lock, insufficient system resources existed to initialize the read-write lock. |
| EDEADLK | The calling thread already has an exclusive lock for <i>rwl</i> .                                                                    |

EFAULT      A fault occurred when the kernel tried to access *rwl*.  
EINVAL      The read-write lock *rwl* is invalid.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_trywrlock()*,  
*pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*

**Synopsis:**

```
#include <pthread.h>
#include <time.h>

int pthread_rwlock_timedrdlock(
 pthread_rwlock_t * rwlock,
 const struct timespec * abs_timeout);
```

**Arguments:**

|                    |                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------|
| <i>rwlock</i>      | The read-write lock that you want to lock.                                                                                 |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> that specifies the maximum time to wait to acquire the lock, expressed as an absolute time. |

**Library:**

`libc`

**Description:**

The *pthread\_rwlock\_timedrdlock()* function applies a read lock to the read-write lock referenced by *rwlock* as in *pthread\_rwlock\_rdlock()*.

However, if the lock can't be acquired without waiting for other threads to unlock it, this wait terminates when the specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the clock on which timeouts are based (i.e. when the value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

The timeout is based on the `CLOCK_REALTIME` clock.

If the read-write lock can be locked immediately, the validity of the *abs\_timeout* parameter isn't checked, and the function won't fail with a timeout.

If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-write lock via a call to *pthread\_rwlock\_timedrdlock()*, upon return from the signal handler the thread resumes waiting for the lock as if it hadn't been interrupted.

The calling thread may deadlock if at the time the call is made it holds a write lock on *rwlock*. The results are undefined if this function is called with an uninitialized read-write lock.

**Returns:**

Zero if the lock for reading on the read-write lock object referenced by *rwlock* is acquired, or an error number to indicate the error.

**Errors:**

|           |                                                                                                                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN    | Couldn't acquire read lock because the maximum number of read locks for lock would be exceeded.                                                                                                      |
| EDEADLK   | The calling thread already holds a write lock on <i>rwlock</i> .                                                                                                                                     |
| EINVAL    | The value specified by <i>rwlock</i> doesn't refer to an initialized read-write lock object, or the <i>abs_timeout</i> nanosecond value is less than zero or greater than or equal to 1,000 million. |
| ETIMEDOUT | The lock couldn't be acquired before the specified timeout expired.                                                                                                                                  |

**Classification:**

POSIX 1003.1j

**Safety**

---

Cancellation point    Yes

Interrupt handler    No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_timedwrlock()*, *pthread\_rwlock\_trywrlock()*,  
*pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_unlock()*,  
*pthread\_rwlock\_wrlock()*, **timespec**

## ***pthread\_rwlock\_timedwrlock()***

© 2004, QNX Software Systems Ltd.

*Lock a read-write lock for writing*

### **Synopsis:**

```
#include <pthread.h>
#include <time.h>

int pthread_rwlock_timedwrlock(
 pthread_rwlock_t * rwlock,
 const struct timespec * abs_timeout);
```

### **Arguments:**

|                    |                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------|
| <i>rwlock</i>      | The read-write lock that you want to lock.                                                                                 |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> that specifies the maximum time to wait to acquire the lock, expressed as an absolute time. |

### **Library:**

**libc**

### **Description:**

The *pthread\_rwlock\_timedwrlock()* function applies a write lock to the read-write lock referenced by *rwlock* as in *pthread\_rwlock\_wrlock()*.

However, if the lock can't be acquired without waiting for other threads to unlock the lock, this wait terminates when the specified timeout expires. The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the clock on which timeouts are based (i.e. when the value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time of the call.

The timeout is based on the **CLOCK\_REALTIME** clock.

If the read-write lock can be locked immediately, the validity of the *abs\_timeout* parameter isn't checked, and the function won't fail with a timeout.

If a signal that causes a signal handler to be executed is delivered to a thread blocked on a read-write lock via a call to *pthread\_rwlock\_timedwrlock()*, upon return from the signal handler the thread resumes waiting for the lock as if it hadn't been interrupted.

The calling thread may deadlock if at the time the call is made it holds a write lock on *rwlock*. The results are undefined if this function is called with an uninitialized read-write lock.

### Returns:

Zero if the lock for writing on the read-write lock object referenced by *rwlock* is acquired, or an error number to indicate the error.

### Errors:

|           |                                                                                                                                                                                                      |
|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN    | Couldn't acquire read lock because the maximum number of read locks for lock would be exceeded.                                                                                                      |
| EDEADLK   | The calling thread already holds the <i>rwlock</i> .                                                                                                                                                 |
| EINVAL    | The value specified by <i>rwlock</i> doesn't refer to an initialized read-write lock object, or the <i>abs_timeout</i> nanosecond value is less than zero or greater than or equal to 1,000 million. |
| ETIMEDOUT | The lock couldn't be acquired before the specified timeout expired.                                                                                                                                  |

### Classification:

POSIX 1003.1j

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_timedrdlock()*, *pthread\_rwlock\_trywrlock()*,  
*pthread\_rwlock\_tryrdlock()*, *pthread\_rwlock\_unlock()*,  
*pthread\_rwlock\_wrlock()*, **timespec**



## ***pthread\_rwlock\_tryrdlock()*** *Attempt to acquire a shared lock on a read-write lock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_tryrdlock(
 pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to lock for reading.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_tryrdlock()` function attempts to acquire a shared lock on the read-write lock referenced by *rwl*. If the read-write lock is already exclusively locked by any thread (including the calling thread), the function returns immediately instead of blocking until a read lock can be obtained.

### **Returns:**

|         |                                                                                                                                        |
|---------|----------------------------------------------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                                                               |
| EAGAIN  | On the first use of a statically initialized read-write lock, insufficient system resources existed to initialize the read-write lock. |
| EBUSY   | The read-write lock was already write locked.                                                                                          |
| EDEADLK | The calling thread already has an exclusive lock for <i>rwl</i> .                                                                      |
| EFAULT  | A fault occurred when the kernel tried to access <i>rwl</i> .                                                                          |
| EINVAL  | The read-write lock <i>rwl</i> is invalid.                                                                                             |

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlock\_destroy(), pthread\_rwlock\_init(),  
pthread\_rwlock\_rdlock(), pthread\_rwlock\_trywrlock(),  
pthread\_rwlock\_unlock(), pthread\_rwlock\_wrlock()*

## ***pthread\_rwlock\_trywrlock()***

*Attempt to acquire an exclusive lock on a read-write lock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_trywrlock(
 pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to lock for writing.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_trywrlock()` function attempts to acquire an exclusive lock on the read-write lock referenced by *rwl*. If the read-write lock is already exclusively locked or shared locked, the function returns immediately instead of blocking until an exclusive lock can be obtained.

The function may need to block to determine the state of the read-write lock.

### **Returns:**

|         |                                                                                                                                        |
|---------|----------------------------------------------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                                                               |
| EAGAIN  | On the first use of a statically initialized read-write lock, insufficient system resources existed to initialize the read-write lock. |
| EBUSY   | The read-write lock was already write locked or read locked.                                                                           |
| EDEADLK | The calling thread already has an exclusive lock for <i>rwl</i> .                                                                      |

EFAULT      A fault occurred when the kernel tried to access *rwl*.  
EINVAL      The read-write lock *rwl* is invalid.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_tryrdlock()*,  
*pthread\_rwlock\_unlock()*, *pthread\_rwlock\_wrlock()*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_unlock(pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to unlock.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_unlock()` function unlocks a read-write lock referenced by *rwl*. The read-write lock may become available for any threads that were blocked on the read-write lock, depending on whether the read-write lock had been locked in exclusive or shared mode.



---

The read-write lock should be owned by the calling thread. If the calling thread doesn't hold the lock, no error status is returned, and the behavior of this read-write lock is now undefined.

---

### **Returns:**

|        |                                                                                                                                        |
|--------|----------------------------------------------------------------------------------------------------------------------------------------|
| EOK    | Success.                                                                                                                               |
| EAGAIN | On the first use of a statically initialized read-write lock, insufficient system resources existed to initialize the read-write lock. |
| EFAULT | A fault occurred when the kernel tried to access <i>rwl</i> .                                                                          |
| EINVAL | The read-write lock <i>rwl</i> is invalid.                                                                                             |

EPERM No thread has a read or write lock on *rwl* or the calling thread doesn't have a write lock on *rwl*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_tryrdlock()*,  
*pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_wrlock()*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlock_wrlock(
 pthread_rwlock_t* rwl);
```

### **Arguments:**

*rwl* A pointer to a `pthread_rwlock_t` object that you want to lock for writing.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlock_wrlock()` function acquires an exclusive lock on the read-write lock referenced by *rwl*. If the read-write lock is already shared-locked by any thread (including the calling thread) or exclusively-locked by any thread (other than the calling thread), the calling thread blocks until all shared locks and exclusive locks are released.

If a signal is delivered to a thread waiting to lock a read-write lock, it resumes waiting for the lock after returning from the signal handler.

### **Returns:**

|         |                                                                                                                                        |
|---------|----------------------------------------------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                                                               |
| EAGAIN  | On the first use of a statically initialized read-write lock, insufficient system resources existed to initialize the read-write lock. |
| EDEADLK | The calling thread already has an exclusive lock for <i>rwl</i> .                                                                      |
| EFAULT  | A fault occurred when the kernel tried to access <i>rwl</i> .                                                                          |

EINVAL      The read-write lock *rwl* is invalid.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlock\_destroy()*, *pthread\_rwlock\_init()*,  
*pthread\_rwlock\_rdlock()*, *pthread\_rwlock\_tryrdlock()*,  
*pthread\_rwlock\_trywrlock()*, *pthread\_rwlock\_unlock()*



## ***pthread\_rwlockattr\_destroy()***

*Destroy a read-write lock attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlockattr_destroy(
 pthread_rwlockattr_t* attr);
```

### **Arguments:**

*attr* A pointer to the `pthread_rwlockattr_t` object that you want to destroy.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlockattr_destroy()` function destroys a read-write lock attribute object created by `pthread_rwlockattr_init()`.



---

Don't use a destroyed read-write lock attribute object reinitializing it by calling `pthread_rwlockattr_init()`.

---

### **Returns:**

EOK Success.  
EINVAL The object specified by *attr* is invalid.

### **Classification:**

Standard Unix

#### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_rwlockattr\_getpshared()*, *pthread\_rwlockattr\_init()*,  
*pthread\_rwlockattr\_setpshared()*

## ***pthread\_rwlockattr\_getpshared()***

*Get the process-shared attribute of a read-write lock attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlockattr_getpshared(
 const pthread_rwlockattr_t* attr,
 int* pshared);
```

### **Arguments:**

*attr*            A pointer to the `pthread_rwlockattr_t` object that you want to get the attribute from.

*pshared*        A pointer to a location where the function can store the process-shared attribute.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlockattr_getpshared()` function gets the the process-shared attribute for the read-write lock attribute object specified by *attr*, storing it in *pshared*.

To let any thread with access to the read-write lock object's memory operate it, the process-shared attribute must be set to `PTHREAD_PROCESS_SHARED`, even if those threads are in different processes. Set the process-shared attribute to `PTHREAD_PROCESS_PRIVATE` to limit access to threads in the current process.

### **Returns:**

`EOK`            Success.

`EINVAL`        The read-write lock attribute object specified by *attr* is invalid.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlockattr\_destroy()*, *pthread\_rwlockattr\_init()*,  
*pthread\_rwlockattr\_setpshared()*

**Synopsis:**

```
#include <pthread.h>

int pthread_rwlockattr_init(
 pthread_rwlockattr_t* attr);
```

**Arguments:**

*attr* A pointer to the `pthread_rwlockattr_t` object that you want to initialize.

**Library:**

`libc`

**Description:**

The `pthread_rwlockattr_init()` function initializes the specified read-write lock attribute object to its default values.

Changes made to a read-write lock attribute object changes after it's been used to initialize a read-write lock won't affect the previously initialized read-write locks.

**Returns:**

EOK Success.  
ENOMEM There isn't enough memory available to initialize *attr*.

**Classification:**

Standard Unix

**Safety**

---

Cancellation point No  
*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_rwlockattr\_destroy()*, *pthread\_rwlockattr\_getpshared()*,  
*pthread\_rwlockattr\_setpshared()*

## ***pthread\_rwlockattr\_setpshared()***

*Set the process-shared attribute of a read-write lock attribute object*

### **Synopsis:**

```
#include <pthread.h>

int pthread_rwlockattr_setpshared(
 pthread_rwlockattr_t* attr,
 int pshared);
```

### **Arguments:**

- attr*            A pointer to the `pthread_rwlockattr_t` object that you want to set the attribute for.
- pshared*        The new value of the process-shared attribute; one of:
- `PTHREAD_PROCESS_SHARED` — let any thread with access to the read-write lock object's memory operate it, even if those threads are in different processes.
  - `PTHREAD_PROCESS_PRIVATE` — limit access to threads in the current process.

### **Library:**

`libc`

### **Description:**

The `pthread_rwlockattr_setpshared()` function sets the process-shared attribute for the read-write lock attribute object specified by *attr* to *pshared*.

### **Returns:**

- `EOK`            Success.
- `EINVAL`        The *pshared* argument is invalid.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_rwlockattr\_destroy()*, *pthread\_rwlockattr\_getpshared()*,  
*pthread\_rwlockattr\_init()*



**Synopsis:**

```
#include <pthread.h>

pthread_t pthread_self(void);
```

**Library:**

libc

**Description:**

The *pthread\_self()* function returns the thread ID of the calling thread.

**Returns:**

The ID of the calling thread.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_equal()*

## ***pthread\_setcancelstate()***

© 2004, QNX Software Systems Ltd.

*Set a thread's cancellation state*

### **Synopsis:**

```
#include <pthread.h>

int pthread_setcancelstate(int state,
 int* oldstate);
```

### **Arguments:**

*state*        The new cancellation state.

*oldstate*     A pointer to a location where the function can store the old cancellation state.

### **Library:**

`libc`

### **Description:**

The *pthread\_setcancelstate()* function sets the calling thread's cancellation state to *state* and returns the previous cancellation state in *oldstate*.

The cancellation state can have the following values:

`PTHREAD_CANCEL_DISABLE`

Cancellation requests are held pending.

`PTHREAD_CANCEL_ENABLE`

Cancellation requests may be acted on according to the cancellation type; see *pthread\_setcanceltype()*.

The default cancellation state for a thread is `PTHREAD_CANCEL_ENABLE`.



---

You can set this attribute (in a non-POSIX way) before creating the thread; for more information, see “QNX extensions,” in the documentation for *pthread\_create()*.

---

### Returns:

EOK          Success.  
EINVAL      The cancellation state specified by *state* is invalid.

### Classification:

POSIX 1003.1 (Threads)

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_cancel()*, *pthread\_setcanceltype()*, *pthread\_testcancel()*

## ***pthread\_setcanceltype()***

© 2004, QNX Software Systems Ltd.

*Set a thread's cancellation type*

### **Synopsis:**

```
#include <pthread.h>

int pthread_setcanceltype(int type,
 int* oldtype);
```

### **Arguments:**

*type*            The new cancellation type.

*oldtype*        A pointer to a location where the function can store the old cancellation type.

### **Library:**

`libc`

### **Description:**

The *pthread\_setcanceltype()* function sets the calling thread's cancellation type to *type* and returns the previous cancellation type in *oldtype*.

The cancellation type can have the following values:

`PTHREAD_CANCEL_ASYNCHRONOUS`

If cancellation is enabled, new or pending cancellation requests may be acted on immediately.

`PTHREAD_CANCEL_DEFERRED`

If cancellation is enabled, cancellation requests are held pending until a cancellation point.

The default cancellation state for a thread is `PTHREAD_CANCEL_DEFERRED`. Note that the standard POSIX and C library calls aren't asynchronous-cancellation safe.



---

You can set this attribute (in a non-POSIX way) before creating the thread; for more information, see “QNX extensions,” in the documentation for *pthread\_create()*.

---

### Returns:

EOK          Success.  
EINVAL      Invalid cancelability type *type*.

### Classification:

POSIX 1003.1 (Threads)

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_cancel()*, *pthread\_setcancelstate()*, *pthread\_testcancel()*

## ***pthread\_setconcurrency()***

© 2004, QNX Software Systems Ltd.

*Set the concurrency level for a thread*

### **Synopsis:**

```
#include <pthread.h>

int pthread_setconcurrency(int new_Level);
```

### **Arguments:**

*new\_Level*     The new value for the concurrency level.

### **Library:**

libc

### **Description:**

QNX Neutrino doesn't support the multiplexing of user threads on top of several kernel scheduled entities. As such, the *pthread\_setconcurrency()* and *pthread\_getconcurrency()* functions are provided for source code compatibility but they have no effect when called. To maintain the function semantics, the *new\_Level* parameter is saved when *pthread\_setconcurrency()* is called so that a subsequent call to *pthread\_getconcurrency()* returns the same value.

### **Returns:**

|        |                                                                                       |
|--------|---------------------------------------------------------------------------------------|
| EOK    | Success.                                                                              |
| EINVAL | Negative argument <i>new_Level</i> .                                                  |
| EAGAIN | The value specified by <i>new_Level</i> would cause a system resource to be exceeded. |

### **Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_getconcurrency()*

## ***pthread\_setschedparam()***

© 2004, QNX Software Systems Ltd.

*Set thread scheduling parameters*

### **Synopsis:**

```
#include <pthread.h>

int pthread_setschedparam(
 pthread_t thread,
 int policy,
 const struct sched_param *param);
```

### **Arguments:**

- thread*     The ID of the thread that you want to get the scheduling parameters for. You can get a thread ID by calling *pthread\_create()* or *pthread\_self()*.
- policy*     The new scheduling policy; one of:
- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority, ready thread runs until it blocks or is preempted by a higher priority thread.
  - SCHED\_RR — the same as SCHED\_FIFO, except threads at the same priority level time slice (round robin) every 50 msec.
  - SCHED\_OTHER — currently the same as SCHED\_RR.



---

Currently, you can set a thread's scheduling policy to SCHED\_SPORADIC only when you create the thread. If you use sporadic scheduling, you can't change the policy later. For more information, see *pthread\_attr\_setschedpolicy()*.

---

*param*     A pointer to a **sched\_param** structure that specifies the scheduling parameters that you want to use.

### **Library:**

**libc**



**Description:**

The *pthread\_setschedparam()* function sets the scheduling policy and associated scheduling parameters of thread *thread* to the values specified in *policy* and *param*.

**Returns:**

|         |                                                                                               |
|---------|-----------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                      |
| EINVAL  | Invalid scheduling policy <i>policy</i> or parameters <i>param</i> .                          |
| ENOTSUP | Unsupported scheduling policy <i>policy</i> or parameters <i>param</i> .                      |
| EPERM   | Insufficient privilege to modify scheduling policy <i>policy</i> or parameters <i>param</i> . |
| ESRCH   | Invalid thread ID <i>thread</i> .                                                             |

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_getschedparam()*, **sched\_param**

## ***pthread\_setspecific()***

© 2004, QNX Software Systems Ltd.

*Set a thread-specific data value*

### **Synopsis:**

```
#include <pthread.h>

int pthread_setspecific(pthread_key_t key,
 const void* value);
```

### **Arguments:**

*key*        The key associated with the data that you want to set. See *pthread\_key\_create()*.

*value*     The value that you want to store.

### **Library:**

libc

### **Description:**

The *pthread\_setspecific()* function binds the thread specific data value *value* with the thread specific data key *key*.

You can call this function from within a thread-specific data destructor function.



---

You must call this function with a key that you got from *pthread\_key\_create()*. You can't use a key after destroying it with *pthread\_key\_delete()*.

---

### **Returns:**

EOK        Success.

ENOMEM    Insufficient memory to store thread specific data value *value*.

EINVAL    Invalid thread specific data key *key*.

## Examples:

See *pthread\_key\_create()*.

## Classification:

POSIX 1003.1 (Threads)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

Calling *pthread\_setspecific()* with a non-NULL *value* may result in lost storage or infinite loops unless *value* was returned by *pthread\_key\_create()*.

## See also:

*pthread\_key\_create()*, *pthread\_getspecific()*

## ***pthread\_sigmask()***

© 2004, QNX Software Systems Ltd.

*Examine and change blocked signals*

### **Synopsis:**

```
#include <signal.h>

int pthread_sigmask(int how,
 const sigset_t* set,
 sigset_t* oset);
```

### **Arguments:**

*how* How you want to change the signal mask; one of:

- SIG\_BLOCK — make the resulting signal mask the union of the current signal mask and the signal set *set*.
- SIG\_UNBLOCK — make the resulting signal mask the intersection of the current signal mask and the complement of the signal set *set*.
- SIG\_SETMASK — make the resulting signal mask the signal set *set*.

This argument is valid only if *set* is non-NULL.

*set* A pointer to a **sigset\_t** object that specifies the signals that you want to affect in the mask.

*oset* NULL, or a pointer to a **sigset\_t** object where the function can store the thread's old signal mask.

### **Library:**

**libc**

### **Description:**

The *pthread\_sigmask()* function is used to examine and/or change the calling thread's signal mask. If *set* is non-NULL, the thread's signal mask is set to *set*. If *oset* is non-NULL, the thread's old signal mask is returned in *oset*.

You can't block the SIGKILL and SIGSTOP signals.

**Returns:**

EOK          Success.  
EINVAL      Invalid *how* parameter.

**Classification:**

POSIX 1003.1 (Threads)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sigprocmask()*

## ***pthread\_sleepon\_broadcast()***

© 2004, QNX Software Systems Ltd.

*Unblock waiting threads*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_broadcast(const volatile void * addr);
```

### **Arguments:**

*addr* The handle that the threads are waiting on. The value of *addr* is typically a data structure that controls a resource.

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_broadcast()* function unblocks all threads currently waiting on *addr*. The threads are unblocked in priority order.

Here's a table to help you decide when to use *pthread\_sleepon\_broadcast()* or *pthread\_sleepon\_signal()*:

| <b>Task</b>                                | <b><i>pthread_sleepon_broadcast()</i></b>                                                                                                                                                              | <b><i>pthread_sleepon_signal()</i></b>                                                                                                                  |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| Mapping a single predicate to one address. | You must recheck the predicate and reblock if necessary. See <i>pthread_sleepon_signal()</i> for a better implementation. The first thread to wake up owns the lock, all others must go back to sleep. | This is the correct and efficient use of <i>pthread_sleepon_signal()</i> . You don't have to recheck the predicate. One thread owns the lock at a time. |

*continued...*

| <b>Task</b>                                 | <b><i>pthread_sleepon_broadcast(pthread_sleepon_signal())</i></b>                                                                                                                                               |                                                                                        |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| Mapping multiple predicates to one address. | The <i>pthread_sleepon_broadcast()</i> function is necessary to wake up all blocked threads. You must recheck the predicates and reblock if necessary. You should try to map only one predicate to one address. | Don't use <i>pthread_sleepon_signal()</i> in this case; it could result in a deadlock. |

**Returns:**

|        |                          |
|--------|--------------------------|
| EOK    | Success.                 |
| EINVAL | Invalid sleepon address. |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_broadcast()*, *pthread\_sleepon\_signal()*,  
*pthread\_sleepon\_lock()*, *pthread\_sleepon\_unlock()*,  
*pthread\_sleepon\_wait()*

## ***pthread\_sleepon\_lock()***

© 2004, QNX Software Systems Ltd.

*Lock the pthread\_sleepon\* functions*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_lock(void);
```

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_lock()* function calls *pthread\_mutex\_lock()* on a mutex associated with the *pthread\_sleepon\** class of functions. You should call this function before testing conditions that determine whether you need to call *pthread\_sleepon\_wait()*, *pthread\_sleepon\_signal()*, or *pthread\_sleepon\_broadcast()*. This mutex prevents other threads from changing the conditions between the time you examine and act upon them.

This function may be implemented as a simple macro.

### **Returns:**

|         |                                                                                           |
|---------|-------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                  |
| EDEADLK | The calling thread already owns the controlling mutex.                                    |
| EAGAIN  | On the first use of <i>pthread_sleepon_lock()</i> , all kernel mutex objects were in use. |

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point Yes

*continued...*



**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_mutex\_lock()*, *pthread\_sleepon\_broadcast()*,  
*pthread\_sleepon\_signal()*, *pthread\_sleepon\_unlock()*,  
*pthread\_sleepon\_wait()*

## ***pthread\_sleepon\_signal()***

© 2004, QNX Software Systems Ltd.

*Signal a sleeping thread*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_signal(const volatile void * addr);
```

### **Arguments:**

*addr* The handle that the threads are waiting on. The value of *addr* is typically a data structure that controls a resource.

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_signal()* function unblocks the highest priority thread waiting on *addr*.

Here's a table to help you decide when to use *pthread\_sleepon\_broadcast()* or *pthread\_sleepon\_signal()*:

| <b>Task</b>                                | <b><i>pthread_sleepon_broadcast()</i></b>                                                                                                                                                              | <b><i>pthread_sleepon_signal()</i></b>                                                                                                         |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Mapping a single predicate to one address. | You must recheck the predicate and reblock if necessary. See <i>pthread_sleepon_signal()</i> for a better implementation. The first thread to wake up owns the lock, all others must go back to sleep. | This is the correct and efficient use of <i>pthread_sleepon_signal()</i> . You must recheck the predicate. One thread owns the lock at a time. |

*continued...*

| <b>Task</b>                                 | <i>pthread_sleepon_broadcast(pthread_sleepon_signal())</i>                                                                                                                                                      |                                                                                        |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| Mapping multiple predicates to one address. | The <i>pthread_sleepon_broadcast()</i> function is necessary to wake up all blocked threads. You must recheck the predicates and reblock if necessary. You should try to map only one predicate to one address. | Don't use <i>pthread_sleepon_signal()</i> in this case; it could result in a deadlock. |

**Returns:**

|        |                          |
|--------|--------------------------|
| EOK    | Success.                 |
| EINVAL | Invalid sleepon address. |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_signal()*, *pthread\_sleepon\_broadcast()*,  
*pthread\_sleepon\_lock()*, *pthread\_sleepon\_unlock()*,  
*pthread\_sleepon\_wait()*

## ***pthread\_sleepon\_timedwait()***

© 2004, QNX Software Systems Ltd.

*Make a thread sleep while waiting*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_timedwait(const volatile void * addr,
 uint64_t nsec);
```

### **Arguments:**

*addr*     The handle that you want the thread to wait for. The value of *addr* is typically a data structure that controls a resource.

*nsec*     A limit on the amount of time to wait, in nanoseconds.

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_timedwait()* function uses a mutex and a condition variable to sleep on a handle, *addr*.

If *nsec* is nonzero, then *pthread\_sleepon\_timedwait()* calls *pthread\_cond\_timedwait()*. If the *pthread\_cond\_timedwait()* times out, then *pthread\_sleepon\_timedwait()* returns ETIMEDOUT. If *nsec* is zero, then *pthread\_sleepon\_timedwait()* calls *pthread\_cond\_wait()* instead.

The *pthread\_sleepon\*()* functions provide a simple, uniform way to wait on a variety of resources in a multithreaded application. For example, a multithreaded filesystem may wish to wait on such diverse things as a cache block, a file lock, an operation complete and many others. For example, to wait on a resource:

```
pthread_sleepon_lock();

while((ptr = cachelist->free) == NULL) {
 pthread_sleepon_timedwait(cachelist);
}
cachelist->free = ptr->free;

pthread_sleepon_unlock();
```

To start an operation and wait upon its completion:

```
/* Line up for access to the driver */
pthread_sleepon_lock();
if(driver->busy) {
 pthread_sleepon_timedwait(&driver->busy);
}

/* We now have exclusive use of the driver */
driver->busy = 1;
driver_start(driver); /* This should be relatively fast */

/* Wait for something to signal driver complete */
pthread_sleepon_timedwait(&driver->complete);
pthread_sleepon_unlock();

/* Get the status/data */
driver_complete(driver);

/* Release control of the driver and signal anyone waiting */
pthread_sleepon_lock();
driver->busy = 0;
pthread_sleepon_signal(&driver->busy);
pthread_sleepon_unlock();
```

The use of a `while` loop instead of an `if` handles the case where the wait on *addr* is woken up using *pthread\_sleepon\_broadcast()*.

You must call *pthread\_sleepon\_lock()*, which acquires the controlling mutex for the condition variable and ensures that another thread won't enter the critical section between the test, block and use of the resource. Since *pthread\_sleepon\_timedwait()* calls *pthread\_cond\_timedwait()*, it releases the controlling mutex when it blocks. It reacquires the mutex before waking up.

The wakeup is accomplished by another thread's calling *pthread\_sleepon\_signal()*, which wakes up a single thread, or *pthread\_sleepon\_broadcast()*, which wakes up all threads blocked on *addr*. Threads are woken up in priority order. If there's more than one thread with the same highest priority, the one that has been waiting the longest is woken first.

A single mutex and one condition variable for each unique address that's currently being blocked on are used. The total number of

condition variables is therefore equal to the number of unique *addrs* that have a thread waiting on them. This also means that the maximum number of condition variables never exceeds the number of threads. To accomplish this, condition variables are dynamically created as needed and placed upon an internal freelist for reuse when not.

You might find the *pthread\_sleepon\_\**(*)* functions easier to use and understand than condition variables. They also resemble the traditional *sleepon()* and *wakeup()* functions found in Unix kernels. They can be implemented as follows:

```
int _sleepon(void *addr) {
 int ret;

 if((ret = pthread_sleepon_lock()) == EOK) {
 ret = pthread_sleepon_timedwait(addr);
 pthread_sleepon_unlock();
 }
 return ret;
}

void _wakeup(void *addr) {
 if(pthread_sleepon_lock() == EOK) {
 pthread_sleepon_broadcast(addr);
 pthread_sleepon_unlock();
 }
}
```

Note that in most Unix kernels, a thread runs until it blocks and thus need not worry about protecting the condition it checks with a mutex. Likewise when a Unix *wakeup()* is called, there isn't an immediate thread switch. Therefore, you can use only the above simple routines (*\_wakeup()* and *\_sleepon()*) if all your threads run with SCHED\_FIFO scheduling and at the same priority, thus more closely mimicking Unix kernel scheduling.

### Returns:

- |           |                                                        |
|-----------|--------------------------------------------------------|
| EDEADLK   | The calling thread already owns the controlling mutex. |
| ETIMEDOUT | The time specified by <i>nsec</i> has passed.          |

EOK                      Success.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_cond\_wait(), pthread\_mutex\_lock(), pthread\_mutex\_unlock(),  
pthread\_sleepon\_broadcast(), pthread\_sleepon\_lock(),  
pthread\_sleepon\_signal(), pthread\_sleepon\_unlock(),  
pthread\_sleepon\_wait(), sched\_setscheduler()*

## ***pthread\_sleepon\_unlock()***

© 2004, QNX Software Systems Ltd.

*Unlock the pthread\_sleepon\*() functions*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_unlock(void);
```

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_unlock()* function calls *pthread\_mutex\_unlock()* on a mutex associated with the *pthread\_sleepon\*()* class of functions. You should call it at the end of a critical section entered by *pthread\_sleepon\_lock()*.

This function may be implemented as a simple macro.

### **Returns:**

EOK        Success.  
EPERM     The current thread doesn't own the controlling mutex.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*pthread\_mutex\_unlock(), pthread\_sleepon\_broadcast(),  
pthread\_sleepon\_lock(), pthread\_sleepon\_signal(),  
pthread\_sleepon\_wait()*

## ***pthread\_sleepon\_wait()***

© 2004, QNX Software Systems Ltd.

*Make a thread sleep while waiting*

### **Synopsis:**

```
#include <pthread.h>

int pthread_sleepon_wait(const volatile void * addr);
```

### **Arguments:**

*addr*     The handle that you want the thread to wait for. The value of *addr* is typically a data structure that controls a resource.

### **Library:**

libc

### **Description:**

The *pthread\_sleepon\_wait()* function uses a mutex and a condition variable to sleep on a handle, *addr*.

The *pthread\_sleepon\** functions provide a simple, uniform way to wait on a variety of resources in a multithreaded application. For example, a multithreaded filesystem may wish to wait on such diverse things as a cache block, a file lock, an operation complete and many others. For example, to wait on a resource:

```
pthread_sleepon_lock();

while((ptr = cachelist->free) == NULL) {
 pthread_sleepon_wait(cachelist);
}
cachelist->free = ptr->free;

pthread_sleepon_unlock();
```

To start an operation and wait for its completion:

```
/* Line up for access to the driver */
pthread_sleepon_lock();
if(driver->busy) {
 pthread_sleepon_wait(&driver->busy);
}
```

```
/* We now have exclusive use of the driver */
driver->busy = 1;
driver_start(driver); /* This should be relatively fast */

/* Wait for something to signal driver complete */
pthread_sleepon_wait(&driver->complete);
pthread_sleepon_unlock();

/* Get the status/data */
driver_complete(driver);

/* Release control of the driver and signal anyone waiting */
pthread_sleepon_lock();
driver->busy = 0;
pthread_sleepon_signal(&driver->busy);
pthread_sleepon_unlock();

pthread_exit(NULL);
```

Choose carefully when you decide whether to use a **while** loop

- If the wait on *addr* is woken up using *pthread\_sleepon\_broadcast()*, you must use a **while** loop.
- If threads are woken up using *pthread\_sleepon\_signal()*, you may use the **if** conditional if the design of the program guarantees proper synchronization and scheduling among contending threads. This is guaranteed in the above example, assuming that none of the threads attempt to reacquire the driver resource (i.e. *pthread\_exit()* call).

If you're in doubt, use a **while** loop, because it guarantees access to the desired resource.

You must call *pthread\_sleepon\_lock()*, which acquires the controlling mutex for the condition variable and ensures that another thread won't enter the critical section between the test, block and use of the resource. Since *pthread\_sleepon\_wait()* calls *pthread\_cond\_wait()*, it releases the controlling mutex when it blocks. It reacquires the mutex before waking up.

The wakeup is accomplished by another thread's calling *pthread\_sleepon\_signal()*, which wakes up a single thread, or

*pthread\_sleepon\_broadcast()*, which wakes up all threads blocked on *addr*. Threads are woken up in priority order. If there's more than one thread with the same highest priority, the one that has been waiting the longest is woken first.

A single mutex and one condition variable for each unique address that's currently being blocked on are used. The total number of condition variables is therefore equal to the number of unique *addrs* that have a thread waiting on them. This also means that the maximum number of condition variables never exceeds the number of threads. To accomplish this, condition variables are dynamically created as needed and placed upon an internal freelist for reuse when not.

You might find the *pthread\_sleepon\_\**(*)* functions easier to use and understand than condition variables. They also resemble the traditional *sleepon()* and *wakeup()* functions found in Unix kernels. They can be implemented as follows:

```
int _sleepon(void *addr) {
 int ret;

 if((ret = pthread_sleepon_lock()) == EOK) {
 ret = pthread_sleepon_wait(addr);
 pthread_sleepon_unlock();
 }
 return ret;
}

void _wakeup(void *addr) {
 if(pthread_sleepon_lock() == EOK) {
 pthread_sleepon_broadcast(addr);
 pthread_sleepon_unlock();
 }
}
```

Note that in most Unix kernels, a thread runs until it blocks, and thus need not worry about protecting the condition it checks with a mutex. Likewise, when a Unix *wakeup()* is called, there isn't an immediate thread switch. Therefore, you can use only the above simple routines (*\_wakeup()* and *\_sleepon()*) if all your threads run with SCHED\_FIFO scheduling and at the same priority, thus more closely mimicking Unix kernel scheduling.

## Returns:

|         |                                                        |
|---------|--------------------------------------------------------|
| EOK     | Success.                                               |
| EDEADLK | The calling thread already owns the controlling mutex. |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_cond\_wait()*, *pthread\_mutex\_lock()*, *pthread\_mutex\_unlock()*,  
*pthread\_sleepon\_broadcast()*, *pthread\_sleepon\_lock()*,  
*pthread\_sleepon\_signal()*, *pthread\_sleepon\_unlock()*,  
*sched\_setscheduler()*

# ***pthread\_spin\_destroy()***

© 2004, QNX Software Systems Ltd.

*Destroy a thread spinlock*

## **Synopsis:**

```
#include <pthread.h>

int pthread_spin_destroy(
 pthread_spinlock_t * spinner);
```

## **Arguments:**

*spinner*     A pointer to the `pthread_spinlock_t` object that you want to destroy.

## **Library:**

`libc`

## **Description:**

The `pthread_spin_destroy()` function destroys the thread spinlock *spinner*, releasing its resources.

Once you've destroyed the spinlock, don't use it again until you've reinitialized it by calling `pthread_spin_init()`.

Calling `pthread_spin_destroy()` gives undefined results when a thread has *spinner* locked or when *spinner* isn't initialized.

## **Returns:**

|        |                                                                                        |
|--------|----------------------------------------------------------------------------------------|
| EOK    | Success.                                                                               |
| EBUSY  | The thread spinlock <i>spinner</i> is in use by another thread and can't be destroyed. |
| EINVAL | Invalid <code>pthread_spinlock_t</code> object <i>spinner</i> .                        |

## **Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_spin\_init()*, *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*,  
*pthread\_spin\_unlock()*

## ***pthread\_spin\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize a thread spinlock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_spin_init(pthread_spinlock_t * spinner,
 int pshared);
```

### **Arguments:**

- spinner* A pointer to the `pthread_spinlock_t` object that you want to initialize.
- pshared* The value that you want to use for the process-shared attribute of the spinlock. The possible values are:
- `PTHREAD_PROCESS_SHARED` — the spinlock may be operated on by any thread that has access to the memory where the spinlock is allocated, even if it's allocated in memory that's shared by multiple processes.
  - `PTHREAD_PROCESS_PRIVATE` — the spinlock can be operated on only by threads created within the same process as the thread that initialized the spinlock. If threads of differing processes attempt to operate on such a spinlock, the behavior is undefined.

### **Library:**

`libc`

### **Description:**

The `pthread_spin_init()` function allocates the resources required for the thread spinlock `spinner`, and initializes `spinner` to an unlocked state.

Any thread that can access the memory where `spinner` is allocated can operate on the spinlock.



Results are undefined if you call *pthread\_spin\_init()* on a *spinner* that's already initialized, or if you try to use a spinlock that hasn't been initialized.

**Returns:**

Zero on success, or an error number to indicate the error.

**Errors:**

|        |                                                                                              |
|--------|----------------------------------------------------------------------------------------------|
| EAGAIN | The system doesn't have the resources required to initialize a new spinlock.                 |
| EBUSY  | The process spinlock, <i>spinner</i> , is in use by another thread and can't be initialized. |
| EINVAL | Invalid <code>pthread_spinlock_t</code> object <i>spinner</i> .                              |
| ENOMEM | The system doesn't have enough free memory to create the new spinlock.                       |

**Classification:**

POSIX 1003.1j (draft)

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_spin\_destroy()*, *pthread\_spin\_lock()*, *pthread\_spin\_trylock()*, *pthread\_spin\_unlock()*

## ***pthread\_spin\_lock()***

© 2004, QNX Software Systems Ltd.

*Lock a thread spinlock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_spin_lock(pthread_spinlock_t * spinner);
```

### **Arguments:**

*spinner*      A pointer to the `pthread_spinlock_t` object that you want to lock.

### **Library:**

`libc`

### **Description:**

The `pthread_spin_lock()` function locks the thread spinlock specified by *spinner*. If *spinner* isn't immediately available, `pthread_spin_lock()` blocks until *spinner* can be locked.

If a thread attempts to lock a spinlock that's already locked via `pthread_spin_lock()` or `pthread_spin_trylock()`, the thread returns EDEADLK.

### **Returns:**

|         |                                                                 |
|---------|-----------------------------------------------------------------|
| EOK     | Success.                                                        |
| EAGAIN  | Insufficient resources available to lock <i>spinner</i> .       |
| EDEADLK | The calling thread already holds <i>spinners</i> lock.          |
| EINVAL  | Invalid <code>pthread_spinlock_t</code> object <i>spinner</i> . |

### **Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

You may not get the desired behavior from this function because the spinlocks are implemented using mutexes. If you really want to use a spinlock, you must use the *pthread\_spin\_lock()* function.

**See also:**

*pthread\_spin\_destroy()*, *pthread\_spin\_init()*, *pthread\_spin\_trylock()*,  
*pthread\_spin\_unlock()*

## ***pthread\_spin\_trylock()***

© 2004, QNX Software Systems Ltd.

*Try to lock a thread spinlock*

### **Synopsis:**

```
#include <pthread.h>

int pthread_spin_trylock(
 pthread_spinlock_t * spinner);
```

### **Arguments:**

*spinner* A pointer to the `pthread_spinlock_t` object that you want to try to lock.

### **Library:**

`libc`

### **Description:**

The `pthread_spin_trylock()` function attempts to lock the thread spinlock specified by *spinner*. It returns immediately if *spinner* can't be locked.

If a thread attempts to lock a spinlock that it's already locked via `pthread_spin_lock()` or `pthread_spin_trylock()`, the thread deadlocks.

### **Returns:**

|        |                                                                         |
|--------|-------------------------------------------------------------------------|
| EOK    | Success.                                                                |
| EAGAIN | Insufficient resources available to lock <i>spinner</i> .               |
| EBUSY  | The thread spinlock <i>spinner</i> is already locked by another thread. |
| EINVAL | Invalid <code>pthread_spinlock_t</code> object <i>spinner</i> .         |

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_spin\_destroy()*, *pthread\_spin\_init()*, *pthread\_spin\_lock()*,  
*pthread\_spin\_unlock()*

# ***pthread\_spin\_unlock()***

© 2004, QNX Software Systems Ltd.

*Unlock a thread spinlock*

## **Synopsis:**

```
#include <pthread.h>

int pthread_spin_unlock(spinlock_t * spinner);
```

## **Arguments:**

*spinner* A pointer to the `pthread_spinlock_t` object that you want to unlock.

## **Library:**

`libc`

## **Description:**

The `pthread_spin_unlock()` function unlocks the thread spinlock specified by *spinner*, which was locked with `pthread_spin_lock()` or `pthread_spin_trylock()`.

If there are threads spinning on *spinner*, the spinlock becomes available, and an unspecified thread acquires the lock.

## **Returns:**

|        |                                           |
|--------|-------------------------------------------|
| EOK    | Success.                                  |
| EINVAL | Invalid process spinlock <i>spinner</i> . |
| EPERM  | The calling thread doesn't hold the lock. |

## **Classification:**

POSIX 1003.1j (draft)

### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*pthread\_spin\_destroy()*, *pthread\_spin\_init()*, *pthread\_spin\_lock()*,  
*pthread\_spin\_trylock()*

# ***pthread\_testcancel()***

© 2004, QNX Software Systems Ltd.

*Test thread cancellation*

## **Synopsis:**

```
#include <pthread.h>

void pthread_testcancel(void);
```

## **Library:**

libc

## **Description:**

The *pthread\_testcancel()* function creates a cancellation point in the calling thread. This function has no effect if cancellation is disabled.

## **Classification:**

POSIX 1003.1 (Threads)

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*pthread\_cancel()*, *pthread\_setcancelstate()*, *pthread\_setcanceltype()*, *ThreadCancel()*



### **Synopsis:**

```
#include <pthread.h>

int pthread_timedjoin(
 pthread_t thread,
 void** value_ptr,
 const struct timespec* abstime);
```

### **Arguments:**

*thread*        The target thread whose termination you're waiting for.

*value\_ptr*    NULL, or a pointer to a location where the function can store the value passed to *pthread\_exit()* by the target thread.

*abstime*      A pointer to a **timespec** structure that specifies the maximum time to wait for the join, expressed as an absolute time.

### **Library:**

**libc**

### **Description:**

The *pthread\_timedjoin()* function is similar to *pthread\_join()*, except that an error of ETIMEDOUT is returned if the join doesn't occur before the absolute time specified by *abstime* passes (i.e. the system time is greater than or equal to *abstime*):

If you are not too long, I will wait here for you all my life.

— Oscar Wilde, *The Importance of Being Earnest*

The *pthread\_timedjoin()* function blocks the calling thread until the target thread *thread* terminates, unless *thread* has already terminated.

If *value\_ptr* is non-NULL and *pthread\_timedjoin()* returns successfully, then the value passed to *pthread\_exit()* by the target thread is placed in *value\_ptr*. If the target thread has been canceled then *value\_ptr* is set to PTHREAD\_CANCELED.

The target thread must be joinable. Multiple *pthread\_join()*, *pthread\_timedjoin()*, *ThreadJoin()*, and *ThreadJoin\_r()* calls on the same target thread aren't allowed. When *pthread\_timedjoin()* returns successfully, the target thread has been terminated.

**Returns:**

|           |                                                                                |
|-----------|--------------------------------------------------------------------------------|
| EOK       | Success.                                                                       |
| EBUSY     | The thread <i>thread</i> is already being joined.                              |
| EDEADLK   | The thread <i>thread</i> is the calling thread.                                |
| EFAULT    | A fault occurred trying to access the buffers provided.                        |
| EINTR     | The function call was interrupted.                                             |
| EINVAL    | The thread <i>thread</i> isn't joinable.                                       |
| ESRCH     | The thread <i>thread</i> doesn't exist.                                        |
| ETIMEDOUT | The absolute time specified in <i>abstime</i> passed before the join occurred. |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_create()*, *pthread\_detach()*, *pthread\_exit()*, *pthread\_join()*,  
*ThreadJoin()*, *ThreadJoin\_r()*, **timespec**

## **\_pulse**

© 2004, QNX Software Systems Ltd.

*Structure that describes a pulse*

### **Synopsis:**

```
#include <sys/neutrino.h>

struct _pulse {
 uint16_t type;
 uint16_t subtype;
 int8_t code;
 uint8_t zero[3];
 union signalval value;
 int32_t scoid;
};
```

### **Description:**

The `_pulse` structure describes a *pulse*, a fixed-size, nonblocking message that carries a small payload (four bytes of data plus a single byte code). The members include:

*type*        `_PULSE_TYPE` (0)

*subtype*    `_PULSE_SUBTYPE` (0)

*code*        A code that identifies the type of pulse. The QNX Neutrino OS reserves the negative codes, including the following:

- `_PULSE_CODE_UNBLOCK`
- `_PULSE_CODE_DISCONNECT`
- `_PULSE_CODE_THREADDEATH`
- `_PULSE_CODE_COIDDEATH`
- `_PULSE_CODE_NET_ACK`,  
`_PULSE_CODE_NET_UNBLOCK`, and  
`_PULSE_CODE_NET_DETACH` — reserved for the  
`io_net` resource manager.

You can define your own pulses, with a code in the range from `_PULSE_CODE_MINAVAIL` through `_PULSE_CODE_MAXAVAIL`.

- value* Information that's relevant to the code:
- `_PULSE_CODE_UNBLOCK` — the receive ID (*rcvid*) associated with the blocking message.
  - `_PULSE_CODE_DISCONNECT` — no value defined.
  - `_PULSE_CODE_THREADDEATH` — the thread ID of the thread that died.
  - `_PULSE_CODE_COIDDEATH` — the connection ID of a connection that was attached to a destroyed channel.

For more details, see *ChannelCreate()*.

If you define your own pulses, you can decide what information you want to store in this field.

*scoid* Server connection ID.

## Classification:

QNX Neutrino

## See also:

*ChannelCreate()*, *MsgReceive()*, *MsgReceivePulse()*, *MsgReceivePulsev()*, *MsgReceivev()*, *MsgSendPulse()*, **sigevent**

## ***pulse\_attach()***

© 2004, QNX Software Systems Ltd.

*Attach a handler function to a pulse code*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int pulse_attach(dispatch_t * dpp,
 int flags,
 int code,
 int (* func)
 (message_context_t * ctp,
 int code,
 unsigned flags,
 void * handle),
 void * handle);
```

### **Arguments:**

- dpp*      The dispatch handle, as returned by *dispatch\_create()*.
- flags*     Currently, the following flag is defined in `<sys/dispatch.h>`:
- `MSG_FLAG_ALLOC_PULSE` — allocate and attach a pulse code that's different than any other code that was either given to *pulse\_attach()* through the *code* argument, or allocated by *pulse\_attach()*. The allocated code is in the range `_PULSE_CODE_MINAVAIL` and `_PULSE_CODE_MAXAVAIL`.
- code*      The pulse code that you want to attach the function to.
- func*      The function that you want to call when a message in the given range is received; see “Handler function,” in the documentation for *message\_attach()*.
- handle*    An arbitrary handle that you want to associate with data for the defined message range. This handle is passed to *func*.

**Library:**`libc`**Description:**

The *pulse\_attach()* function attaches a pulse *code* to a user-supplied function *func*. You can use the same function *func* with *message\_attach()*.

When the resource manager receives a pulse that matches *code*, it calls *func*. This user-supplied function is responsible for doing any specific work needed to handle the pulse pointed to by *ctp->msg.pulse*. The *handle* passed to the function is the *handle* initially passed to *pulse\_attach()*. The *handle* may be a device entry you want associated with the pulse *code*.

You typically use *pulse\_attach()* to associate pulses generated by interrupt handlers or timers with a routine in the main program of your resource manager. By examining *ctp->rvid*, the *func* function can determine whether a pulse or message was received.

**Returns:**

If `MSG_FLAG_ALLOC_PULSE` is specified, the function returns the allocated pulse code; otherwise, it returns the *code* that's passed in. On failure, -1 is returned (*errno* is set).

**Errors:**

|        |                                                                    |
|--------|--------------------------------------------------------------------|
| EAGAIN | Couldn't allocate a pulse <i>code</i> .                            |
| EINVAL | The pulse <i>code</i> is out of range, or it's already registered. |
| ENOMEM | Insufficient memory to allocate internal data structures.          |

## Examples:

```
#include <sys/dispatch.h>
#include <time.h>
#include <stdio.h>
#include <stdlib.h>

int my_func(...) {
 :
}

int main(int argc, char **argv) {
 dispatch_t *dpp;
 int flag = 0, code, mycode;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 :

 mycode = ...;

 if ((code = pulse_attach(dpp, flag, mycode,
 &my_func, NULL)) == -1) {
 fprintf (stderr, "Failed to attach code %d.\n", mycode);
 return 1;
 }
 /* else successfully attached a pulse code */

 :
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*message\_attach()*, *pulse\_detach()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

## ***pulse\_detach()***

© 2004, QNX Software Systems Ltd.

*Detach a handler function from a pulse code*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int pulse_detach(dispatch_t * dpp,
 int code,
 int flags);
```

### **Arguments:**

*dpp*      The dispatch handle, as returned by *dispatch\_create()*.

*code*     The pulse code that you want to detach.

*flags*    Reserved; pass 0 for this argument.

### **Library:**

libc

### **Description:**

The *pulse\_detach()* function detaches the pulse *code*, for dispatch handle *dpp*, that was attached with *pulse\_attach()*.

### **Returns:**

0          Success.

-1         The pulse *code* doesn't match any attached pulse code.

### **Examples:**

```
#include <sys/dispatch.h>
#include <time.h>
#include <stdio.h>
#include <stdlib.h>

int my_func(...) {
 :
}
```

```

}

int main(int argc, char **argv) {
 dispatch_t *dpp;
 int flag=0, code, mycode;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 :

 if ((code = pulse_attach(dpp, flag, mycode,
 &my_func, NULL)) == -1) {
 fprintf (stderr, "Failed to attach pulse code %d.\n", \
 mycode);
 return 1;
 }

 :

 if (pulse_detach (dpp, code, flag) == -1) {
 fprintf (stderr, "Failed to detach code %d.\n", code);
 return 1;
 }
 /* else message was detached */

 :
}

```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*message\_detach()*, *pulse\_attach()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

## Synopsis:

```
#include <stdio.h>

int putc(int c,
 FILE* fp);
```

## Arguments:

*c* The character that you want to write.

*fp* The stream you want to write the character on.

## Library:

libc

## Description:

The *putc()* macro writes the character *c*, cast as **(int) (unsigned char)**, to the output stream designated by *fp*.

## Returns:

The character written, cast as **(int) (unsigned char)**, or EOF if an error occurs (*errno* is set).

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 FILE* fp;
 int c;

 fp = fopen("file", "r");
 if(fp != NULL) {
 while((c = fgetc(fp)) != EOF) {
 putc(c, stdout);
 }
 fclose(fp);
 }
}
```

```
 return EXIT_SUCCESS;
 }
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

*putc()* is a macro.

**See also:**

*errno*, *ferror()*, *fopen()*, *fputc()*, *fputchar()*, *fputs()*, *getc()*,  
*getc\_unlocked()*, *getchar()*, *getchar\_unlocked()*, *putchar()*,  
*putchar\_unlocked()*, *putc\_unlocked()*, *puts()*

**Synopsis:**

```
#include <stdio.h>

int putc_unlocked(int c,
 FILE *stream);
```

**Arguments:**

*c*            The character that you want to write.

*stream*       The stream you want to write the character on.

**Library:**

`libc`

**Description:**

The *putc\_unlocked()* function is a thread-unsafe version of *putc()*. You can use it safely only when the invoking thread has locked *stream* using *flockfile()* (or *ftrylockfile()*) and *funlockfile()*.

**Returns:**

The character written, cast as `(int) (unsigned char)`, or EOF if an error occurred (*errno* is set).

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*getc()*, *getchar()*, *getchar\_unlocked()*, *getc\_unlocked()*, *putc()*,  
*putchar()*, *putchar\_unlocked()*



## Synopsis:

```
#include <stdio.h>

int putchar(int c);
```

## Arguments:

*c* The character that you want to write.

## Library:

libc

## Description:

The *putchar()* function writes the character *c*, cast as (int) (unsigned char), to the *stdout* stream. It's equivalent to:

```
fputc(c, stdout);
```

## Returns:

The character written, cast as (int) (unsigned char), or EOF if an error occurs (*errno* is set).

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 FILE *fp;
 int c;

 fp = fopen("file", "r");
 c = fgetc(fp);
 while(c != EOF) {
 putchar(c);
 c = fgetc(fp);
 }
}
```

```
 fclose(fp);
 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, fputc(), fputchar(), fputs(), getc(), getc\_unlocked(), getchar(),  
getchar\_unlocked(), putc(), putchar\_unlocked(), putc\_unlocked()*

### **Synopsis:**

```
#include <stdio.h>

int putchar_unlocked(int c);
```

### **Arguments:**

*c* The character that you want to write.

### **Library:**

`libc`

### **Description:**

The *putchar\_unlocked()* function is a thread-unsafe version of *putchar()*. You can use it safely only when the invoking thread has locked *stdout* using *flockfile()* (or *ftrylockfile()*) and *funlockfile()*.

### **Returns:**

The character written, cast as `(int) (unsigned char)`, or EOF if an error occurred (*errno* is set).

### **Classification:**

POSIX 1003.1

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*getc()*, *getc\_unlocked()*, *getchar()*, *getchar\_unlocked()*, *putc()*,  
*putc\_unlocked()*, *putchar()*

## Synopsis:

```
#include <stdlib.h>

int putenv(const char *env_name);
```

## Arguments:

*env\_name*      The name of the environment, and what you want to do to it; see below.

## Library:

libc

## Description:

The *putenv()* function uses *env\_name*, in the form *name=value*, to set the environment variable *name* to *value*. This function alters *name* if it exists, or creates a new environment variable.

In either case, *env\_name* becomes part of the environment; subsequent modifications to the string pointed to by *env\_name* affect the environment.

The space for environment names and their values is limited. Consequently, *putenv()* can fail when there's insufficient space remaining to store an additional value.



---

If *env\_name* isn't a literal string, you should duplicate the string, since *putenv()* doesn't copy the value. For example:

```
putenv(strdup(buffer));
```

---

## Returns:

- 0 Success.
- 1 An error occurred; *errno* is set.

## Errors:

- ENOMEM There wasn't enough memory to expand the environment.

## Examples:

The following gets the string currently assigned to **INCLUDE** and displays it, assigns a new value to it, gets and displays it, and then removes **INCLUDE** from the environment.

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *path;
 path = getenv("INCLUDE");
 if(path != NULL) {
 printf("INCLUDE=%s\n", path);
 }

 if(putenv("INCLUDE=/src/include") != 0) {
 printf("putenv() failed setting INCLUDE\n");
 return EXIT_FAILURE;
 }

 path = getenv("INCLUDE");
 if(path != NULL) {
 printf("INCLUDE=%s\n", path);
 }

 unsetenv("INCLUDE");

 return EXIT_SUCCESS;
}
```

This program produces the following output:

```
INCLUDE=/usr/nto/include
INCLUDE=/src/include
```

## Classification:

Standard Unix

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## Caveats:

Never use *putenv()* with an automatic variable.

The *putenv()* function manipulates the environment pointed to by the global *environ* variable.

## See also:

*clearenv()*, *environ*, *errno*, *getenv()*, *setenv()*, *unsetenv()*

## ***puts()***

© 2004, QNX Software Systems Ltd.

*Write a string to stdout*

### **Synopsis:**

```
#include <stdio.h>

int puts(const char *buf);
```

### **Arguments:**

*buf* A pointer to the zero-terminated string that you want to write.

### **Library:**

libc

### **Description:**

The *puts()* function writes the character string pointed to by *buf* to the *stdout* stream, and appends a newline character to the output. The terminating NUL character of *buf* isn't written.

### **Returns:**

A nonnegative value for success, or EOF if an error occurs (*errno* is set).

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 FILE *fp;
 char buffer[80];

 fp = freopen("file", "r", stdin);
 while(gets(buffer) != NULL) {
 puts(buffer);
 }
 fclose(fp);

 return EXIT_SUCCESS;
}
```



## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno, fputs(), gets(), putc()*

## ***putspent()***

© 2004, QNX Software Systems Ltd.

*Put an entry into the shadow password database*

### **Synopsis:**

```
#include <sys/types.h>
#include <shadow.h>

int putspent(const struct spwd* p,
 FILE* fp);
```

### **Arguments:**

- p* A pointer to a **spwd** structure that contains the entry that you want to write.
- fp* The stream that you want to write the entry on.

### **Library:**

libc

### **Description:**

The *putspent()* function writes a shadow password entry into the specified file. This function is the inverse of *getspent()*.

Given a pointer to a **spwd** structure created by the *getspent()* or the *getspnam()* routine, *putspent()* writes a line on the stream *fp*, which matches the format of `</etc/shadow>`. The **spwd** structure contains the following members:

```
char *sp_namp; /* name */
char *sp_pwdp; /* encrypted password */
long sp_lstchg; /* last changed */
long sp_max; /* #days (min) to change */
long sp_min; /* #days (max) to change */
long sp_warn; /* #days to warn */
long sp_inact; /* #days of inactivity */
long sp_expire; /* date to auto-expire */
long sp_flag; /* reserved */
```

If the *sp\_min*, *sp\_max*, *sp\_lstchg*, *sp\_warn*, *sp\_inact*, or *sp\_expire* field of the structure is -1, or if *sp\_flag* = 0, the corresponding `</etc/shadow>` field is cleared.

**Returns:**

Zero.

**Errors:**

The *putspent()* function uses the following functions, and as a result *errno* can be set to an error for any of these calls:

- *fclose()*
- *fgets()*
- *fopen()*
- *fseek()*
- *rewind()*

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <pwd.h>
#include <shadow.h>

/*
 * This program adds a user and password to
 * a temporary file which can then be used with
 * fgetspent() (of course the password
 * string should be encrypted already etc.)
 */

int main(int argc, char** argv)
{
 FILE* fp;
 struct spwd sp;
 char pwbuf[80], nambuf[80];

 memset(&sp, 0, sizeof(sp));
 if (argc < 2) {
 printf("%s filename \n", argv[0]);
 return(EXIT_FAILURE);
 }

 if (!(fp = fopen(argv[1], "w"))) {
 fprintf(stderr, "Can't open file %s \n", argv[1]);
 perror("Problem ");
 }
}
```

```
 return(1);
 }

 printf("Enter a userid: ");
 if (!gets(nambuf)) {
 fprintf(stderr, "Can't get username string\n");
 }
 sp.sp_namp = nambuf;

 printf("Enter a password: ");
 if (!gets(pwbuf)) {
 fprintf(stderr, "Can't get username password\n");
 }
 sp.sp_pwdp = pwbuf;

 putspent(&sp, fp);
 fclose(fp);
 return(EXIT_SUCCESS);
}
```

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*errno*, *getgrent()*, *getlogin()*, *getpwnam()*, *getpwuid()*, *getspent()*,  
*getspnam()*, *setspent()*

## Synopsis:

```
#include <utmp.h>

void pututline(struct utmp * __utmp);
```

## Arguments:

*\_\_utmp* A pointer to the **utmp** structure for the entry that you want to add.

## Library:

**libc**

## Description:

The *pututline()* function writes out the supplied **utmp** structure into the **utmp** file.

It uses *getutid()* to search forward for the proper place if it finds that it isn't already there. Normally, you should search for the proper entry by calling *getutent()*, *getutid()*, or *getutline()*. If so, *pututline()* doesn't search. If *pututline()* doesn't find a matching slot for the new entry, it adds a new entry to the end of the file.

When called by a non-**root** user, *pututline()* invokes a *setuid()* root program to verify and write the entry, since the file specified in *\_PATH\_UTMP* is normally writable only by root. In this event, the *ut\_name* field must correspond to the actual user name associated with the process; the *ut\_type* field must be either **USER\_PROCESS** or **DEAD\_PROCESS**; the *ut\_line* field must be a device-special file and be writable by the user.

## Returns:

A pointer to the **utmp** structure.

**Files:**

*\_PATH\_UTMP*

Specifies the user information file.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

The most current entry is saved in a static structure. Copy it before making further accesses.

On each call to either *getutid()* or *getutline()*, the routine examines the static structure before performing more I/O. If the contents of the static structure match what it's searching for, the function looks no further. For this reason, to use *getutline()* to search for multiple occurrences, zero out the static area after each success, or *getutline()* will return the same structure over and over again.

There's one exception to the rule about emptying the structure before further reads are done: the implicit read done by *pututline()* (if it finds that it isn't already at the correct place in the file) doesn't hurt the contents of the static structure returned by the *getutent()*, *getutid()* or *getutline()* routines, if the user has just modified those contents and passed the pointer back to *pututline()*.

These routines use buffered standard I/O for input, but *pututline()* uses an unbuffered nonstandard write to avoid race conditions between processes trying to modify the `utmp` and `wtmp` files.

**See also:**

*endutent()*, *getutent()*, *getutid()*, *getutline()*, *setutent()*, `utmp`,  
*utmpname()*

`login` in the *Utilities Reference*

## ***putw()***

© 2004, QNX Software Systems Ltd.

*Put a word on a stream*

---

### **Synopsis:**

```
#include <wchar.h>

int putw(int w,
 FILE *stream);
```

### **Arguments:**

*w*            The word that you want to write.

*stream*      The stream that you want to write a word on.

### **Library:**

`libc`

### **Description:**

The *putw()* function writes the C `int` (word) *w* to the standard I/O output *stream* (at the position of the file pointer, if defined). The size of a word is the size of an integer, and varies from machine to machine. The *putw()* function neither assumes nor causes special alignment in the file.

### **Returns:**

0      Success.

1      An error occurred; *errno* is set.

### **Errors:**

EFBIG      The file is a regular file and an attempt was made to write at or beyond the offset maximum.



## Classification:

Legacy Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Because of possible differences in word length and byte ordering, files written using *putw()* are machine-dependent, and might not be read correctly using *getw()* on a different processor.

## See also:

*errno*, *ferror()*, *fopen()*, *fputc()*, *fputchar()*, *fputs()*, *getw()*, *putchar()*, *putchar\_unlocked()*, *putc\_unlocked()*, *puts()*

## ***putwc()***

© 2004, QNX Software Systems Ltd.

*Write a wide character to a stream*

### **Synopsis:**

```
#include <wchar.h>

wint_t putwc(wchar_t wc,
 FILE * fp);
```

### **Arguments:**

*wc*     The wide character that you want to write.

*fp*     The stream that you want to write the wide character on.

### **Library:**

`libc`

### **Description:**

The *putwc()* functions writes the wide character specified by *wc*, cast as `(wint_t) (wchar_t)`, to the output stream specified by *fp*.

### **Returns:**

The wide character written, cast as `(wint_t) (wchar_t)`, or WEOF if an error occurs (*errno* is set).



---

If *wc* exceeds the valid wide-character range, the value returned is the wide character written, not *wc*.

---

### **Errors:**

EAGAIN     The O\_NONBLOCK flag is set for *fp* and would have been blocked by this operation.

EBADF     The file descriptor for *fp* isn't valid for writing.

EFBIG     The file exceeds the maximum file size, the process's file size limit, or the function can't write at or beyond the offset maximum.

|       |                                                                                                                                                                                                                                                                                          |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EINTR | A signal terminated the write operation; no data was transferred.                                                                                                                                                                                                                        |
| EIO   | A physical I/O error has occurred or all of the following conditions were met: <ul style="list-style-type: none"><li>• The process is in the background.</li><li>• TOSTOP is set.</li><li>• The process is blocking/ignoring SIGTTOU.</li><li>• The process group is orphaned.</li></ul> |
| EPIPE | Can't write to pipe or FIFO because it's closed; a SIGPIPE signal is also sent to the thread.                                                                                                                                                                                            |

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*errno*, *getwc()*, *getwchar()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

## ***putwchar()***

© 2004, QNX Software Systems Ltd.

*Write a wide character to a stdout*

### **Synopsis:**

```
#include <wchar.h>

wint_t putwchar(wchar_t wc);
```

### **Arguments:**

*wc* The wide character that you want to write.

### **Library:**

`libc`

### **Description:**

The *putwchar()* function writes the wide character *wc*, cast as `(wint_t) (wchar_t)`, to the *stdout* stream. It's equivalent to:

```
fputc(wc, stdout);
```

### **Returns:**

The wide character written, cast as `(wint_t) (wchar_t)` or `WEOF` if an error occurs (*errno* is set).



---

If *wc* exceeds the valid wide-character range, the value returned is the wide character written, not *wc*.

---

### **Errors:**

|        |                                                                                                                                     |
|--------|-------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The <code>O_NONBLOCK</code> flag is set for <i>fp</i> and would have been blocked by this operation.                                |
| EBADF  | The file descriptor for <i>fp</i> isn't valid for writing.                                                                          |
| EFBIG  | The file exceeds the maximum file size, the process's file size limit, or the function can't write at or beyond the offset maximum. |

|       |                                                                                                                                                                                                                                                                                          |
|-------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EINTR | A signal terminated the write operation; no data was transferred.                                                                                                                                                                                                                        |
| EIO   | A physical I/O error has occurred or all of the following conditions were met: <ul style="list-style-type: none"><li>• The process is in the background.</li><li>• TOSTOP is set.</li><li>• The process is blocking/ignoring SIGTTOU.</li><li>• The process group is orphaned.</li></ul> |
| EPIPE | Can't write to pipe or FIFO because it's closed; a SIGPIPE signal is also sent to the thread.                                                                                                                                                                                            |

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:***errno, getwc(), getwchar()*

“Stream I/O functions” and “Wide-character functions” in the summary of functions chapter.

## ***pwrite()*, *pwrite64()***

© 2004, QNX Software Systems Ltd.

*Write into a file without changing the file pointer*

### **Synopsis:**

```
#include <unistd.h>

ssize_t pwrite(int filedes,
 const void* buff,
 size_t nbytes,
 off_t offset);

ssize_t pwrite64(int filedes,
 const void* buff,
 size_t nbytes,
 off64_t offset);
```

### **Arguments:**

*filedes*    The file descriptor for the file you want to write in.

*buff*        A pointer to a buffer that contains the data you want to write.

*nbytes*      The number of bytes to write.

*offset*      The desired position inside the file.

### **Library:**

libc

### **Description:**

The *pwrite()* function performs the same action as *write()*, except that it writes into a given position without changing the file pointer.

The *pwrite64()* function is a 64-bit version of *pwrite()*.

### **Returns:**

The number of bytes actually written, or -1 if an error occurred (*errno* is set).

**Errors:**

|        |                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the write operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for writing.                                                                                       |
| EFBIG  | File is too big.                                                                                                                                                         |
| EINTR  | The write operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EIO    | A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.                                                             |
| ENOSPC | There's no free space remaining on the device containing the file.                                                                                                       |
| ENOSYS | The <i>pwrite()</i> function isn't implemented for the filesystem specified by <i>filedes</i> .                                                                          |
| EPIPE  | An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.                               |

**Classification:**

*pwrite()* is standard Unix; *pwrite64()* is for large-file support

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close(), creat(), dup(), dup2(), errno, fcntl(), lseek(), open(), pipe(),  
pread(), read(), readv(), select(), write(), writev()*



## Synopsis:

```
#include <unistd.h>

char* qnx_crypt(const char* key,
 const char* salt);
```

## Arguments:

*key* A NUL-terminated string (normally a password typed by a user).

*salt* A two-character string chosen from the set [a-zA-Z0-9./]. This function doesn't validate the values for *salt*, and values outside this range may cause undefined behavior. This string is used to perturb the algorithm in one of 4096 different ways.

## Library:

libc

## Description:

The *qnx\_crypt()* function performs password encryption. It's a variant of the standard *crypt()* function that uses an encryption similar to, but not compatible with, the Data Encryption Standard (DES) encryption. This function is provided for compatibility with QNX 4.



---

The *qnx\_crypt()* function checks only the first eight characters of *key*.

---

## Returns:

A pointer to the encrypted value, or NULL on failure.

## Examples:

```
#include <unistd.h>

int main(int argc, char **argv) {
 char salt[3];
```

```
char string[20];
char *result;

strcpy(string, "thomasf");
salt[0] = 'a';
salt[1] = 'B';
salt[2] = 0;

result = qnx_crypt(string, salt);
printf("Result is [%s] --> [%s] \n", string, result);

return 0;
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The return value points to static data that's overwritten by each call to *qnx\_crypt()*.

**See also:**

*crypt()*, *encrypt()*, *getpass()*, *setkey()*  
**login** in the *Utilities Reference*

## Synopsis:

```
#include <stdlib.h>

void qsort(void* base,
 size_t num,
 size_t width,
 int (*compare) (
 const void* ,
 const void*));
```

## Arguments:

|                |                                                                                                                                                                                                                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>base</i>    | A pointer to the array that you want to sort.                                                                                                                                                                                                                                                      |
| <i>num</i>     | The number of elements in the array.                                                                                                                                                                                                                                                               |
| <i>width</i>   | The size of each element, in bytes.                                                                                                                                                                                                                                                                |
| <i>compare</i> | A pointer to a function that compares two entries. It's called with two arguments that point to elements in the array. The comparison function must return an integer less than, equal to, or greater than zero if the first argument is less than, equal to, or greater than the second argument. |

## Library:

libc

## Description:

The *qsort()* function sorts the *base* array using the comparison function specified by *compare*. The array must have at least *num* elements, each of *width* bytes.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char* some_strs[] = { "last", "middle", "first" };

int compare(const void* op1, const void* op2)
{
 const char **p1 = (const char **) op1;
 const char **p2 = (const char **) op2;

 return(strcmp(*p1, *p2));
}

int main(void)
{
 qsort(some_strs,
 sizeof(some_strs) / sizeof(char *),
 sizeof(char *),
 compare);

 printf("%s %s %s\n",
 some_strs[0], some_strs[1], some_strs[2]);

 return EXIT_SUCCESS;
}
```

produces the output:

```
first last middle
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*bsearch()*



### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

int Raccept(int s,
 struct sockaddr * addr,
 int * addrlen);
```

### **Arguments:**

*s*            A socket that's been created with *socket()*.

*addr*        A result parameter that's filled in with the address of the connecting entity, as known to the communications layer. The exact format of the *addr* parameter is determined by the domain in which the connection was made.

*addrlen*    A value-result parameter. It should initially contain the amount of space pointed to by *addr*; on return it contains the actual length (in bytes) of the address returned. This call is used with connection-based socket types, currently with SOCK\_STREAM.

### **Library:**

**libsocks**

### **Description:**

The *Raccept()* function is a cover function for *accept()* — the difference is that *Raccept()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

## Returns:

A descriptor for the accepted socket, or -1 if an error occurs (*errno* is set).

## Classification:

SOCKS

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*accept()*, *Rbind()*, *Rconnect()*, *Rgetsockname()*, *Rlisten()*, *Rrcmd()*, *Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall



## Synopsis:

```
#include <signal.h>

int raise(int condition);
```

## Arguments:

*condition*    The signal that you want to raise. For more information, see *signal()*.

## Library:

libc

## Description:

The *raise()* function generates the signal specified by *condition*. Use *SignalAction()* or *signal()* to specify the actions to take when a signal is received.

## Returns:

0 if the specified *condition* is sent, or nonzero if an error occurs (*errno* is set).

The *raise()* function doesn't return if the action for that signal is to terminate the program or to transfer control using the *longjmp()* function.

## Errors:

EAGAIN    Insufficient system resources are available to deliver the signal.

EINVAL    The value of *condition* isn't a valid signal number.

## Examples:

Wait until a SIGINT signal is received. The signal is automatically raised on iteration 10000, or when you press Ctrl – C:

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>

sig_atomic_t signal_count;
sig_atomic_t signal_number;

void alarm_handler(int signum)
{
 ++signal_count;
 signal_number = signum;
}

int main(void)
{
 unsigned long i;

 signal_count = 0;
 signal_number = 0;
 signal(SIGINT, alarm_handler);

 printf("Iteration: ");
 for(i = 0; i < 100000; ++i) {
 printf("\b\b\b\b\b\b%d", 5, i);

 if(i == 10000) raise(SIGINT);

 if(signal_count > 0) break;
 }

 if(i == 100000) {
 printf("\nNo signal was raised.\n");
 } else if(i == 10000) {
 printf("\nSignal %d was raised by the "
 "raise() function.\n", signal_number);
 } else {
 printf("\nUser raised signal #%d.\n",
 signal_number);
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*signal()*, *SignalAction()*

## ***rand()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random integer*

---

### **Synopsis:**

```
#include <stdlib.h>

int rand(void);
```

### **Library:**

libc

### **Description:**

The *rand()* function computes a pseudo-random integer in the range 0 to RAND\_MAX. You can start the sequence at different values by calling *srand()*.

The *rand\_r()* function is a thread-safe version of *rand()*.

### **Returns:**

A pseudo-random integer.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int i;

 for(i=1; i < 10; ++i) {
 printf("%d\n", rand());
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | No  |

## See also:

*lrand48(), nrand48(), rand\_r(), srand()*

## ***rand\_r()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random integer in a thread-safe manner*

### **Synopsis:**

```
#include <stdlib.h>

int rand_r(unsigned int* seed);
```

### **Arguments:**

*seed* A pointer to the seed for the sequence of pseudo-random numbers. If you call *rand\_r()* with the same initial value for the *seed*, the same sequence is generated.

### **Library:**

`libc`

### **Description:**

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, *rand\_r()* computes a sequence of pseudo-random integers in the range 0 to `RAND_MAX`.

### **Returns:**

A pseudo-random integer.

### **Classification:**

POSIX 1003.1

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*rand()*, *srand()*

## ***random()***

© 2004, QNX Software Systems Ltd.

*Generate a pseudo-random number from the default state*

### **Synopsis:**

```
#include <stdlib.h>

long random(void);
```

### **Library:**

libc

### **Description:**

The *random()* function uses a nonlinear additive feedback random-number generator employing a default state array size of 31 long integers to return successive pseudo-random numbers in the range from 0 to 231-1. The period of this random-number generator is approximately  $16 \times (231-1)$ . The size of the state array determines the period of the random-number generator. Increasing the state array size increases the period.

Use this function in conjunction with the following:

- initstate()* Initialize the state of the pseudo-random number generator.
- setstate()* Specify the state of the pseudo-random number generator.
- srandom()* Set the seed used by the pseudo-random number generator.

The *random()* and *srandom()* functions have (almost) the same calling sequence and initialization properties as *rand()* and *srand()*. The difference is that *rand()* produces a much less random sequence. In fact, the low dozen bits generated by *rand()* go through a cyclic pattern. All the bits generated by *random()* are usable. For example,

```
random() & 01
```



produces a random binary value.

Unlike *srand()*, *srandom()* doesn't return the old seed because the amount of state information used is much more than a single word. The *initstate()* and *setstate()* routines are provided to deal with restarting/changing random number generators. With 256 bytes of state information, the period of the random-number generator is greater than 269.

Like *rand()*, *random()* produces by default a sequence of numbers that can be duplicated by calling *srandom()* with 1 as the seed.

If *initstate()* hasn't been called, *random()* behaves as though *initstate()* had been called with *seed*=1 and *size*=128.

If *initstate()* is called with *size* less than 8, *random()* uses a simple linear congruential random number generator.

## Returns:

The generated pseudo-random number.

## Examples:

See *initstate()*.

## Classification:

Standard Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | No  |

**See also:**

*drand48(), initstate(), rand(), setstate(), srand(), srandom()*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

int Rbind(int s,
 const struct sockaddr * name,
 int namelen);
```

### **Arguments:**

*s*            The file descriptor to be bound.

*name*        A pointer to the **sockaddr** structure that holds the address to be bound to the socket. The socket length and format depend upon its address family.

*namelen*     The length of the **sockaddr** structure pointed to by *name*.

### **Library:**

**libsocks**

### **Description:**

The *Rbind()* function is a cover function for *bind()* — the difference is that *Rbind()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

### **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

## Classification:

SOCKS

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*bind()*, *Raccept()*, *Rconnect()*, *Rgetsockname()*, *Rlisten()*, *Rrcmd()*,  
*Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall

**Synopsis:**

```
#include <unistd.h>

int rcmd(char ** ahost,
 unsigned short inport,
 const char * locuser,
 const char * remuser,
 const char * cmd,
 int * fd2p);
```

**Arguments:**

|                |                                                                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ahost</i>   | The name of the host that you want to execute the command on. If the function can find the host, it sets * <i>ahost</i> to the standard name of the host. |
| <i>inport</i>  | The well-known Internet port on the host, where the server resides.                                                                                       |
| <i>locuser</i> | The user ID on the local machine.                                                                                                                         |
| <i>remuser</i> | The user ID on the remote machine.                                                                                                                        |
| <i>cmd</i>     | The command that you want to execute.                                                                                                                     |
| <i>fd2p</i>    | See below.                                                                                                                                                |

**Library:**

**libsocket**

**Description:**

The *rcmd()* function is used by the superuser to execute a command, *cmd*, on a remote machine using an authentication scheme based on reserved port numbers. The **rshd** server (among others) uses the *rcmd()*, *rresvport()*, and *ruserok()* functions.

The *rcmd()* function looks up the host \**ahost* by means of *gethostbyname()*, and returns -1 if the host doesn't exist. Otherwise,

\**ahost* is set to the standard name of the host and a connection is established to a server residing at the well-known Internet port *inport*.

If the connection succeeds, a SOCK\_STREAM socket in the Internet domain is returned to the caller and given to the remote command as standard input and standard output.

---

**If *fd2p* is:    Then:**

---

|         |                                                                                                                                                                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nonzero | An auxiliary channel to a control process is set up, and a descriptor for it is placed in * <i>fd2p</i> . The control process will return diagnostic output from the command (unit 2) on this channel and will accept bytes as signal numbers to be forwarded to the command's process group. |
| Zero    | The standard error (unit 2 of the remote command) is made the same as the standard output and no provision is made for sending arbitrary signals to the remote process (although you may be able to get its attention by using out-of-band data).                                             |

The protocol is described in detail in **rshd** in the *Utilities Reference*.

### Returns:

A valid socket descriptor; or -1 if an error occurs and a message is printed to standard error.

### Errors:

The error code EAGAIN is overloaded to mean "All network ports in use."

### Classification:

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*gethostbyname()*, *rresvport()*, *ruserok()*

*rlogin*, *rlogind*, *rsh*, *rshd* in the *Utilities Reference*

## ***Rconnect()***

© 2004, QNX Software Systems Ltd.

*Initiate a connection on a socket (via a SOCKS server)*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

int Rconnect(int s,
 const struct sockaddr * name,
 int namelen);
```

### **Arguments:**

*s*            The descriptor of the socket on which to initiate the connection.

*name*        The name of the socket to connect to for a SOCK\_STREAM connection.

*namelen*     The length of the *name*, in bytes.

### **Library:**

**libsocks**

### **Description:**

The *Rconnect()* function is a cover function for *connect()* — the difference is that *Rconnect()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

### **Returns:**

0            Success.

-1          An error occurred (*errno* is set).



**Classification:**

SOCKS

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*connect()* *Raccept()*, *Rbind()*, *Rgetsockname()*, *Rlisten()*, *Rrcmd()*,  
*Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall

## ***rdchk()***

© 2004, QNX Software Systems Ltd.

*Check to see if a read is likely to succeed*

### **Synopsis:**

```
#include <unix.h>

int rdchk(int fd);
```

### **Arguments:**

*fd*     The file descriptor that you want to check.

### **Library:**

`libc`

### **Description:**

The *rdchk()* function checks to see if a read from the file descriptor, *fd*, is likely to succeed.

### **Returns:**

The number of characters waiting to be read, or -1 if an error occurred (*errno* is set).

### **Errors:**

ENOTTY     The *fd* argument isn't the file descriptor for a character device.

### **Classification:**

Unix

#### **Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*tcischars()*

## ***re\_comp()***

© 2004, QNX Software Systems Ltd.

*Compile a regular expression*

---

### **Synopsis:**

```
#include <unix.h>

char *re_comp(char *s);
```

### **Arguments:**

*s* A string that contains the regular expression that you want to compile. This string must end with a null byte and may include newline characters. If this argument is NULL, the current regular expression remains unchanged.

### **Library:**

libc

### **Description:**

The *re\_comp()* function converts a regular expression string (RE) into an internal form suitable for pattern matching. Use this function with *re\_exec()*.

The *re\_comp()* and *re\_exec()* functions support simple regular expressions. The regular expressions of the form  $\{m\}$ ,  $\{m, \}$ , or  $\{m, n\}$  aren't supported.



For better portability, use *regcomp()*, *regerror()*, and *regexexec()* instead of these functions.

---

### **Returns:**

NULL if the string pointed to by *s* was successfully converted. Otherwise, a pointer to one of the following error message strings is returned:

- **No previous regular expression**
- **Regular expression too long**

- unmatched \ (
- missing ]
- too many \(\) pairs
- unmatched \)

### Classification:

Legacy Unix

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*re\_exec()*, *regcomp()*, *regerror()*, *regexexec()*

**grep** in the *Utilities Reference*

## ***re\_exec()***

© 2004, QNX Software Systems Ltd.

*Execute a regular expression*

### **Synopsis:**

```
#include <unix.h>

int re_exec(char *s);
```

### **Arguments:**

*s* A pointer to the string that you want to compare to the current regular expression. This string must end with a null byte and may include newline characters.

### **Library:**

libc

### **Description:**

The *re\_exec()* function compares the string pointed to by the string argument with the last regular expression passed to *re\_comp()*.

The *re\_comp()* and *re\_exec()* functions support simple regular expressions. The regular expressions of the form  $\{m\}$ ,  $\{m, \}$ , or  $\{m, n\}$  aren't supported.



---

For better portability, use *regcomp()*, *regerror()*, and *regex()* instead of these functions.

---

### **Returns:**

- 1 The string matches the last compiled regular expression.
- 0 The string doesn't match the last compiled regular expression.
- 1 The compiled regular expression is invalid (indicating an internal error).

## Classification:

Legacy Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*re\_comp()*, *regcomp()*, *regerror()*, *regexec()*

**grep** in the *Utilities Reference*

# ***read()***

© 2004, QNX Software Systems Ltd.

*Read bytes from a file*

## **Synopsis:**

```
#include <unistd.h>

ssize_t read(int fildes,
 void* buf,
 size_t nbyte);
```

## **Arguments:**

- fildes*    The descriptor of the file that you want to read from.
- buf*        A pointer to a buffer where the function can store the data that it reads.
- nbyte*     The number of bytes that you want to read.

## **Library:**

`libc`

## **Description:**

The *read()* function attempts to read *nbyte* bytes from the file associated with the open file descriptor, *fildes*, into the buffer pointed to by *buf*.

If *nbyte* is zero, *read()* returns zero, and has no other effect.

On a regular file or other file capable of seeking, *read()* starts at a position in the file given by the file offset associated with *fildes*. Before successfully returning from *read()*, the file offset is incremented by the number of bytes actually read.



---

The *read()* call ignores advisory locks that you may have set with *fcntl()*.

---

On a file not capable of seeking, *read()* starts at the current position.

When *read()* returns successfully, its return value is the number of bytes actually read and placed in the buffer. This number will never



be greater than *nbyte*, although it may be less than *nbyte* for one of the following reasons:

- The number of bytes left in the file is less than *nbyte*.
- The *read()* request was interrupted by a signal.
- The file is a pipe (or FIFO) or a special file, and has fewer than *nbyte* bytes immediately available for reading. For example, reading from a file associated with a terminal may return one typed line of data.

If *read()* is interrupted by a signal before it reads any data, it returns a value of -1 and sets *errno* to EINTR. However, if *read()* is interrupted by a signal after it has successfully read some data, it returns the number of bytes read.

No data is transferred past the current end-of-file. If the starting position is at or after the end-of-file, *read()* returns zero. If the file is a device special file, the result of subsequent calls to *read()* will work, based on the then current state of the device (that is, the end of file is transitory).

If the value of *nbyte* is greater than INT\_MAX, *read()* returns -1 and sets *errno* to EINVAL. See <limits.h>.

When attempting to read from an empty pipe or FIFO:

- 1 If no process has the pipe open for writing, *read()* returns 0 to indicate end-of-file.
- 2 If a process has the pipe open for writing, and O\_NONBLOCK is set, *read()* returns -1, and *errno* is set to EAGAIN.
- 3 If a process has the pipe open for writing, and O\_NONBLOCK is clear, *read()* blocks until some data is written, or the pipe is closed by all processes that had opened it for writing.

When attempting to read from a file (other than a pipe or FIFO) that support nonblocking reads and has no data currently available:

- 1 If O\_NONBLOCK is set, *read()* returns -1, and *errno* is set to EAGAIN.

- 2 If O\_NONBLOCK is clear, *read()* blocks until some data is available.
- 3 The O\_NONBLOCK flag has no effect if some data is available.

If you call *read()* on a portion of a file, prior to the end-of-file, that hasn't been written, it returns bytes with the value zero.

If *read()* succeeds, the *st\_atime* field of the file is marked for update.

**Returns:**

The number of bytes actually read, or -1 (*errno* is set).

**Errors:**

|            |                                                                                                                                                                         |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN     | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the read operation.                                                             |
| EBADF      | The file descriptor, <i>files</i> , isn't a valid file descriptor open for reading.                                                                                     |
| EINTR      | The read operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EIO        | A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.                                                            |
| ENOSYS     | The <i>read()</i> function isn't implemented for the filesystem specified by <i>files</i> .                                                                             |
| E_OVERFLOW | The file is a regular file and an attempt is made to read at or beyond the offset maximum associated with the file.                                                     |

**Examples:**

```
#include <sys/types.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>

int main(void)
{
 int fd;
 int size_read;
 char buffer[80];

 /* Open a file for input */
 fd = open("myfile.dat", O_RDONLY);

 /* Read the text */
 size_read = read(fd, buffer,
 sizeof(buffer));

 /* Test for error */
 if(size_read == -1) {
 perror("Error reading myfile.dat");
 return EXIT_FAILURE;
 }

 /* Close the file */
 close(fd);

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close(), creat(), dup(), dup2(), errno, fcntl(), lseek(), open(), pipe(),  
readblock(), readv(), select(), write(), writev()*

**Synopsis:**

```
#include <snmp/snmp_api.h>

int read_main_config_file(
 struct snmpd_conf_data * info);
```

**Arguments:**

*info* A pointer to a **snmpd\_conf\_data** structure that the function can fill with the configuration information. For more information about this structure, see below.

**Library:**

**libsnmp**

**Description:**

The *read\_main\_config\_file()* function fills the *info* structure with data from the **snmpd.conf** file (see the *Utilities Reference*). This information is useful if you wish to know what configuration information the SNMP agent was started with.

The string pointers in this structure, if not NULL, point to strings obtained by using *malloc()*; you can free them by calling *free()*.



---

If the data for a member of the structure isn't available, the structure member isn't modified. You should use *memset()* to set the structure to 0 before calling *read\_main\_config\_file()*.

---

To locate the **snmpd.conf** file, this function first checks the **SNMPCONFIGFILE** environment variable. If this isn't found, the default, **snmpd.conf**, is used. If the specified file couldn't be accessed, the structure members aren't updated.

The **snmpd\_conf\_data** structure is defined in **<snmp\_api.h>**, and contains the following members:

```
struct snmpd_conf_data{
 char* main_config_fname;
 char* party_conf_fname;
 char* view_conf_fname;
 char* context_conf_fname;
 char* acl_conf_fname;
 char* sysContact;
 char* sysLocation;
 char* sysName;
 char* private_community;
 char* public_community;
 char* trap_sink;
 char* trap_community;
 int conf_authentraps;
};
```

The members of this structure are:

*main\_config\_fname*

**snmpd.conf** file location.

*party\_conf\_fname*

**party.conf** file location.

*view\_conf\_fname*

**view.conf** file location.

*context\_conf\_fname*

**context.conf** file location.

*acl\_conf\_fname*

**acl.conf** file location.

*sysContact* *system.sysContact* string.

*sysLocation* *system.sysLocation* string.

*sysName* *system.sysName* string.

*private\_community*

Private level community string name to use.

*public\_community*

Public level community string name to use.

*trap\_sink*

Destination host to send trap messages.

*trap\_community*

Community name to use for trap messages.

*conf\_authentraps*

Enable authentication traps (1 means enable, 2 means disable).

### Returns:

0 Success.

-1 An error occurred (*errno* is set).

### Errors:

ENOENT The *snmpd.conf* file wasn't found.

### Files:

*snmpd.conf* Default SNMP configuration file. For more information, see the *Utilities Reference*.

### Environment variables:

**SNMPCONFIGFILE**

Location of the SNMP configuration file.

### Classification:

SNMP

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*snmp\_close()*, *snmp\_free\_pdu()*, *snmp\_open()*, **snmp\_pdu**,  
*snmp\_pdu\_create()*, *snmp\_read()*, *snmp\_select\_info()*, *snmp\_send()*,  
**snmp\_session**, *snmp\_timeout()*

**snmpd**, **snmpd.conf** in the *Utilities Reference*



### **Synopsis:**

```
#include <unistd.h>

int readblock(int fd,
 size_t blksize,
 unsigned block,
 int numblks,
 void *buff);
```

### **Arguments:**

|                |                                                                            |
|----------------|----------------------------------------------------------------------------|
| <i>fd</i>      | The file descriptor for the file you want to read from.                    |
| <i>blksize</i> | The number of bytes in each block of data.                                 |
| <i>block</i>   | The block number from which to start reading.                              |
| <i>numblks</i> | The number of blocks to read.                                              |
| <i>buff</i>    | A pointer to a buffer where the function can store the data that it reads. |

### **Library:**

`libc`

### **Description:**

The *readblock()* function reads *numblks* blocks of data from the file associated with the open file descriptor *fd*, into the buffer pointed to by *buff*, starting at block number *block* (blocks are numbered starting at 0). The *blksize* argument specifies the size of a block, in bytes.

This function is useful for direct access to raw blocks on a block special device (for example, raw disk blocks) but may also be used for high-speed access to database files, for example. (The speed gain is through the combined seek/read implicit in this call, and the ability to transfer more than the *read()* function's limit of INT\_MAX bytes at a time.)

If *numblks* is zero, *readblock()* returns zero and has no other results.

On successful completion, *readblock()* returns the number of blocks actually read and placed in the buffer. This number is never greater than *numblks*. The value returned may be less than *numblks* if one of the following occurs:

- The number of blocks left before the end-of-file is less than *numblks*.
- The process requests more blocks than implementation limits allow to be read in a single atomic operation.
- A read error occurred after reading at least one block.

If a read error occurs on the first block, *readblock()* returns -1 and sets *errno* to EIO.

## Returns:

The number of blocks actually read. If an error occurs, it returns -1, sets *errno* to indicate the error, and the contents of the buffer pointer to by *buff* are left unchanged.

## Errors:

|        |                                                                                                |
|--------|------------------------------------------------------------------------------------------------|
| EBADF  | The <i>fd</i> argument isn't a valid file descriptor open for reading a block-oriented device. |
| EIO    | A physical read error occurred on the first block.                                             |
| EINVAL | The starting position is invalid (0 or negative) or beyond the end of the disk.                |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*writeblock()*

## ***readcond()***

© 2004, QNX Software Systems Ltd.

*Read data from a terminal device*

### **Synopsis:**

```
#include <unistd.h>

int readcond(int fd,
 void * buf,
 int n,
 int min,
 int time,
 int timeout);
```

### **Arguments:**

- fd*      The file descriptor associated with the terminal device that you want to read from.
- buf*     A pointer to a buffer into which *readcond()* can put the data.
- n*        The maximum number of bytes to read.
- min, time, timeout*
- When used in RAW mode, these arguments override the behavior of the *MIN* and *TIME* members of the terminal's **termios** structure. For more information, see below.

### **Library:**

**libc**

### **Description:**


The *readcond()* function reads up to *n* bytes of data from the terminal device indicated by *fd* into the buffer pointed to by *buf*.

This function is an alternative to the *read()* function for terminal devices, providing additional arguments for timed read operations. These additional arguments can be used to minimize overhead when dealing with terminal devices.

The three arguments (*min*, *time*, and *timeout*), when used on terminal devices in RAW mode, override the behavior of the *MIN* and *TIME*

elements of the currently defined `termios` structure (for the duration of this call only). The `termios` structure also defines a forwarding character (in `c_cc[VFWDD]`) that can be used to bypass *min*, *time* and *timeout*.

The normal case of a simple read by an application would block until at least one character was available.



|         |                                                          |
|---------|----------------------------------------------------------|
| MIN     | Respond when at least this number of characters arrives. |
| TIME    | Respond if a pause in the character stream occurs.       |
| TIMEOUT | Respond if an overall amount of time passes.             |
| FORWARD | Respond if a framing character arrives.                  |

---

*Conditions that satisfy an input request.*

In the case where multiple conditions are specified, the read will be satisfied when any one of them is satisfied.

### ***MIN***

The qualifier *MIN* is useful when an application has knowledge of the number of characters it expects to receive.

Any protocol that knows the character count for a frame of data can use *MIN* to wait for the entire frame to arrive. This significantly reduces IPC and process scheduling. *MIN* is often used in conjunction with *TIME* or *TIMEOUT*. *MIN* is part of the POSIX standard.

### ***TIME***

The qualifier *TIME* is useful when an application is receiving streaming data and wishes to be notified when the data stops or pauses. The pause time is specified in 1/10 of a second. *TIME* is part of the POSIX standard.

### *TIMEOUT*

The qualifier *TIMEOUT* is useful when an application has knowledge of how long it should wait for data before timing out. The timeout is specified in 1/10 of a second.

Any protocol that knows the character count for a frame of data it expects to receive can use *TIMEOUT*. This in combination with the baud rate allows a reasonable guess to be made when data should be available. It acts as a deadman timer to detect dropped characters. It can also be used in interactive programs with user input to timeout a read if no response is available within a given time.

*TIMEOUT* is a QNX extension and isn't part of the POSIX standard.

### *FORWARD*

The qualifier *FORWARD* is useful when a protocol is delimited by a special framing character. For example, the PPP protocol used for TCP/IP over a serial link start and end its packets with a framing character. When used in conjunction with *TIMEOUT*, the *FORWARD* character can greatly improve the efficiency of a protocol implementation. The protocol process will receive complete frames, rather than character by character. In the case of a dropped framing character, *TIMEOUT* or *TIME* can be used to quickly recover.

This greatly minimizes the amount of IPC work for the OS and results in a much lower processor utilization for a given TCP/IP data rate. It is interesting to note that PPP doesn't contain a character count for its frames. Without the data-forwarding character, an implementation would be forced to read the data one character at a time.

*FORWARD* is a QNX extension and isn't part of the POSIX standard.

To enable the *FORWARD* character, you must set the VFWD character in the *c\_cc* member of the `termios` structure:

```
/* PPP forwarding character */
const char fwd_char = 0x7e;

#include <termios.h>

:
```

```

int fd;
struct termios termio;

:

tcgetattr(fd, &termio);
termio.c_cc[VFWD] = fwd_char;
tcsetattr(fd, TCSANOW, &termio);

```

The following table summarizes the interaction of *min*, *time*, and *timeout*:

| <i>min</i> | <i>time</i> | <i>timeout</i> | Description                                                                                                                                                                                                  |
|------------|-------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 0          | 0           | 0              | Returns immediately with as many bytes as are currently available (up to <i>n</i> bytes).                                                                                                                    |
| <i>M</i>   | 0           | 0              | Return with up to <i>n</i> bytes only when at least <i>M</i> bytes are available.                                                                                                                            |
| 0          | <i>T</i>    | 0              | Return with up to <i>n</i> bytes when at least one byte is available, or $T * .1$ sec has expired.                                                                                                           |
| <i>M</i>   | <i>T</i>    | 0              | Return with up to <i>n</i> bytes when either <i>M</i> bytes are available, or at least one byte has been received and the inter-byte time between any subsequently received characters exceeds $T * .1$ sec. |
| 0          | 0           | <i>t</i>       | RESERVED.                                                                                                                                                                                                    |
| <i>M</i>   | 0           | <i>t</i>       | Return with up to <i>n</i> bytes when $t * .1$ sec has expired, or <i>M</i> bytes are available.                                                                                                             |
| 0          | <i>T</i>    | <i>t</i>       | RESERVED.                                                                                                                                                                                                    |

*continued...*

| <i>min</i> | <i>time</i> | <i>timeout</i> | Description                                                                                                                                                                                                                                                                 |
|------------|-------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>M</i>   | <i>T</i>    | <i>t</i>       | Return with up to <i>n</i> bytes when <i>M</i> bytes are available, or <i>t</i> * .1 sec has expired and no characters are received, or at least one byte has been received and the inter-byte time between any subsequently received characters exceeds <i>T</i> * .1 sec. |

Note that when *timeout* is zero, the behavior of *min* and *time* is exactly the same as the behavior of the **MIN** and **TIME** parameters of the **termios** controlling structure. Thus, *readcond()* can be used as a higher speed alternative to consecutive calls of *tcgetattr()*, *tcsetattr()*, and *read()* when dealing with RAW terminal I/O.

The (*M*, 0, *t*) case is useful for communications protocols that cannot afford to block forever waiting for data that may never arrive.

The (*M*, *T*, *t*) case is provided to permit *readcond()* to return when a burst of data ends (as in the (*M*, *T*, 0) case), but also to return if no burst at all is detected within a reasonable amount of time.

## Returns:

The number of bytes read, or -1 if an error occurs (*errno* is set).

## Errors:

|        |                                                                                                                             |
|--------|-----------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set on this <i>fd</i> , and the process would have been blocked in trying to perform this operation. |
| EBADF  | The argument <i>fd</i> is invalid or file isn't opened for reading.                                                         |
| EINTR  | The <i>readcond()</i> call was interrupted by the process being signalled.                                                  |
| EIO    | This process isn't currently able to read data from this <i>fd</i> .                                                        |



ENOSYS      This function isn't supported for this *fd*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *read()*, *tcgetattr()*, *tcsetattr()*, **termios**

## ***readdir()***

© 2004, QNX Software Systems Ltd.

*Read a directory entry*

---

### **Synopsis:**

```
#include <dirent.h>

struct dirent * readdir(DIR * dirp);
```

### **Arguments:**

*dirp*     A pointer to the directory stream to be read.

### **Library:**

**libc**

### **Description:**

The *readdir()* function reads the next directory entry from the directory specified by *dirp*, which is the value returned by a call to *opendir()*.

You can call *readdir()* repeatedly to list all of the entries contained in the directory specified by the pathname given to *opendir()*. The *closedir()* function must be called to close the directory stream and free the memory allocated by *opendir()*.

The `<dirent.h>` file defines the `struct dirent` and the `DIR` type used by the *readdir()* family of functions.



The result of using a directory stream after one of the *exec\*()* or *spawn\*()* family of functions is undefined. After a call to *fork()*, either the parent *or* the child (but not both) can continue processing the directory stream, using the *readdir()* and *rewinddir()* functions. If both the parent and child processes use these functions, the result is undefined. Either (or both) processes may use *closedir()*.

---

The `<dirent.h>` file also defines the following macros for accessing extra data associated with the `dirent` structure:

*\_DEXTRA\_FIRST( pdirent )*

Get a pointer to the first block of data associated with the structure pointed to by *pdirent*.

*\_DEXTRA\_NEXT( last)*

Get the block of data that follows the block pointed to by *last*.

*\_DEXTRA\_VALID( extra, pdirent)*

Evaluates to 1 if *extra* is a pointer to a valid block of data associated with the structure pointed to by *pdirent*.

You can use these macros to traverse the data associated with the **dirent** structure like this:

```
for(extra = _DEXTRA_FIRST(dirent);
 _DEXTRA_VALID(extra, dirent);
 extra = _DEXTRA_NEXT(extra)) {
 switch(extra->d_type) {
 /* No data */
 case _DTYPE_NONE :
 break;
 /* Data includes information as returned by stat() */
 case _DTYPE_STAT :
 break;
 /* Data includes information as returned by lstat() */
 case _DTYPE_LSTAT :
 break;
 ...
 }
}
```

**Returns:**

A pointer to a **struct dirent** object for success, or NULL if the end of the directory stream is encountered or an error occurs (*errno* is set).

**Errors:**

EBADF            The *dirp* argument doesn't refer to an open directory stream.

EOVERFLOW      One of the values in the structure to be returned can't be represented correctly.

## Examples:

Get a list of files contained in the directory `/home/fred`:

```
#include <stdio.h>
#include <dirent.h>
#include <stdlib.h>

int main(void)
{
 DIR* dirp;
 struct dirent* direntp;

 dirp = opendir("/home/fred");
 if(dirp != NULL) {
 for(;;) {
 direntp = readdir(dirp);
 if(direntp == NULL) break;

 printf("%s\n", direntp->d_name);
 }

 closedir(dirp);

 return EXIT_SUCCESS;
 }

 return EXIT_FAILURE;
}
```

## Classification:

POSIX 1003.1

### Safety

Cancellation point    Yes

Interrupt handler    No

*continued...*

**Safety**

---

|                |    |
|----------------|----|
| Signal handler | No |
| Thread         | No |

**See also:**

*closedir()*, *errno*, *lstat()*, *opendir()*, *readdir\_r()*, *rewinddir()*, *seekdir()*, *telldir()*, *stat()*

## ***readdir\_r()***

© 2004, QNX Software Systems Ltd.

*Get information about the next matching filename*

### **Synopsis:**

```
#include <sys/types.h>
#include <dirent.h>

int readdir_r(DIR * dirp,
 struct dirent * entry,
 struct dirent ** result);
```

### **Arguments:**

*dirp*      A pointer to the directory stream to be read.

*entry*     A pointer to a **dirent** structure where the function can store the directory entry.

*result*    The address of a location where the function can store a pointer to the information found.

### **Library:**

**libc**

### **Description:**

If `_POSIX_THREAD_SAFE_FUNCTIONS` is defined, *readdir\_r()* initializes the **dirent** structure referenced by *entry* with the directory entry at the current position in the directory stream referred to by *dirp*, and stores a pointer to this structure in *result*.

The storage pointed by *entry* must be large enough for a **dirent** structure with the *s\_name* member an array of char containing at least `NAME_MAX` plus one element.



---

The `struct dirent` structure *doesn't* include space for the pathname. You must provide it:

```
struct dirent *entry;
entry = malloc(sizeof(*entry) + NAME_MAX + 1);
```

---

## Returns:

EOK      Success.

On failure, *errno* is set.

## Errors:

EBADF      The *dirp* argument doesn't refer to an open directory stream.

E\_OVERFLOW      One of the values in the structure to be returned can't be represented correctly.

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*closedir(), errno, opendir(), readdir(), seekdir(), telldir(), rewinddir()*



## Synopsis:

```
#include <unistd.h>

int readlink(const char* path,
 char* buf,
 size_t bufsiz);
```

## Arguments:

- path*      The name of the symbolic link.
- buf*        A pointer to a buffer where the function can store the contents of the symbolic link (i.e. the path linked to).
- bufsiz*     The size of the buffer.

## Library:

`libc`

## Description:

The *readlink()* function places the contents of the symbolic link named by *path* into the buffer pointed to by *buf*, which has a size of *bufsiz*. The contents of the returned symbolic link doesn't include a NULL terminator. Its length must be determined from the `stat` structure returned by *lstat()*, or by the return value of the *readlink()* call.

If *readlink()* is successful, up to *bufsiz* bytes from the contents of the symbolic link are placed in *buf*.

## Returns:

The number of bytes placed in the buffer, or -1 if an error occurs (*errno* is set).

## Errors:

|              |                                                                                                                                                    |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of the <i>path</i> prefix.                                                                             |
| EINVAL       | The named file isn't a symbolic link.                                                                                                              |
| ELOOP        | A loop exists in the symbolic links encountered during resolution of the path argument, and more than SYMLOOP_MAX symbolic links were encountered. |
| ENAMETOOLONG | A component of the <i>path</i> exceeded NAME_MAX characters, or the entire pathname exceeded PATH_MAX characters.                                  |
| ENOENT       | The named file doesn't exist.                                                                                                                      |
| ENOSYS       | Links aren't supported by the resource manager associated with path.                                                                               |
| ENOTDIR      | A component of the <i>path</i> prefix named by path isn't a directory.                                                                             |

## Examples:

```
/*
 * Read the contents of the named symbolic links
 */
#include <limits.h>
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

char buf[PATH_MAX + 1];

int main(int argc, char** argv)
{
 int n;
 int len;
 int ecode = 0;

 for(n = 1; n < argc; ++n) {
 if((len = readlink(argv[n], buf, PATH_MAX)) == -1) {
 perror(argv[n]);
 }
 }
}
```

```
 ecode++;
 }
 else {
 buf[len] = '\0';
 printf("%s -> %s\n", argv[n], buf);
 }
}

return(ecode);
}
```

## Classification:

POSIX 1003.1a

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *lstat()*, *symlink()*

# ***readv()***

© 2004, QNX Software Systems Ltd.

*Read bytes from a file*

## **Synopsis:**

```
#include <sys/uio.h>

ssize_t readv(int fdes,
 const iove_t* iov,
 int iovcnt);
```

## **Arguments:**

- fdes*      The descriptor of the file that you want to read from.
- iov*        An array of `iovec_t` objects where the function can store the data that it reads.
- iovcnt*    The number of entries in the *iov* array. The maximum number of entries is `UIO_MAXIOV`.

## **Library:**

`libc`

## **Description:**

The *readv()* function attempts to read from the file associated with the open file descriptor, *fdes*, placing the data into *iovcnt* buffers specified by the members of the *iov* array: *iov*[0], *iov*[1], ..., *iov*[*iovcnt*-1].

On a regular file or other file capable of seeking, *readv()* starts at a position in the file given by the file offset associated with *fdes*. Before successfully returning from *readv()*, the file offset is incremented by the number of bytes actually read.

The *iovec\_t* structure contains the following members:

- iov\_base*      The base address of a memory area where data should be placed.
- iov\_len*        The length of the memory area.

The *readv()* function always fills one buffer completely before proceeding to the next.



---

The *readv()* call ignores advisory locks that may have been set by the *fcntl()* function.

---

On a file not capable of seeking, *readv()* starts at the current position.

When *readv()* returns successfully, its return value is the number of bytes actually read and placed in the buffer. This number will never be greater than the combined sizes of the *iov* buffers, although it may be less for one of the following reasons:

- The number of bytes left in the file is less than the combined size of the *iov* buffers.
- The *readv()* request was interrupted by a signal.
- The file is a pipe (or FIFO) or a special file, and has fewer bytes immediately available for reading. For example, reading from a file associated with a terminal may return one typed line of data.

If *readv()* is interrupted by a signal before it reads any data, it returns a value of -1 and sets *errno* to EINTR. However, if *readv()* is interrupted by a signal after it has successfully read some data, it returns the number of bytes read.

No data is transferred past the current end-of-file. If the starting position is at or after the end-of-file, *readv()* returns zero. If the file is a device special file, the result of subsequent calls to *readv()* will work, based on the then current state of the device (that is, the end of file is transitory).

When attempting to read from an empty pipe or FIFO:

- 1 If no process has the pipe open for writing, *readv()* returns 0 to indicate end-of-file.
- 2 If a process has the pipe open for writing, and O\_NONBLOCK is set, *readv()* returns -1 and sets *errno* to EAGAIN.

- 3 If a process has the pipe open for writing, and O\_NONBLOCK is clear, *read()* blocks until some data is written, or the pipe is closed by all processes that had opened it for writing.

When attempting to read from a file (other than a pipe or FIFO) that supports nonblocking reads and has no data currently available:

- 1 If O\_NONBLOCK is set, *readv()* returns -1 and sets *errno* to EAGAIN.
- 2 If O\_NONBLOCK is clear, *readv()* blocks until some data is available.
- 3 The O\_NONBLOCK flag has no effect if some data is available.

If you call *readv()* on a portion of a file, prior to the end-of-file, that hasn't been written, it returns bytes with the value zero.

If *readv()* succeeds, the *st\_atime* field of the file is marked for update.

## Returns:

The number of bytes read, or -1 if an error occurred (*errno* is set).

## Errors:

|        |                                                                                                                                                                         |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the read operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for reading.                                                                                      |
| EINTR  | The read operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EINVAL | The <i>iovcnt</i> argument is less than or equal to 0, or greater than UIO_MAXIOV.                                                                                      |

|            |                                                                                                                     |
|------------|---------------------------------------------------------------------------------------------------------------------|
| EIO        | A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.        |
| ENOSYS     | The <i>readv()</i> function isn't implemented for the filesystem specified by <i>filedes</i> .                      |
| E_OVERFLOW | The file is a regular file and an attempt is made to read at or beyond the offset maximum associated with the file. |

## Classification:

Standard Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*close()*, *creat()*, *dup()*, *dup2()*, *errno*, *fcntl()*, *lseek()*, *open()*, *pipe()*, *read()*, *readblock()*, *select()*, *write()*, *writenv()*

## ***realloc()***

© 2004, QNX Software Systems Ltd.

*Allocate, reallocate or free a block of memory*

### **Synopsis:**

```
#include <stdlib.h>

void* realloc(void* old_blk,
 size_t size);
```

### **Arguments:**

*old\_blk*     A pointer to the block of memory to be allocated, reallocated, or freed.

*size*        The new size, in bytes, for the block of memory.

### **Library:**

`libc`

### **Description:**

The *realloc()* function allocates, reallocates, or frees the block of memory specified by *old\_blk* based on the following rules:

- If *old\_blk* is NULL, a new block of memory of *size* bytes is allocated.
- If the *size* is zero, the *free()* function is called to release the memory pointed to by *old\_blk*.
- Otherwise, *realloc()* reallocates space for an object of *size* bytes by:
  - shrinking the size of the allocated memory block *old\_blk* when *size* is smaller than the current size of *old\_blk*
  - extending the allocated size of the allocated memory block *old\_blk* if there is a large enough block of unallocated memory immediately following *old\_blk*
  - allocating a new block with the appropriate *size*, and copying the contents of *old\_blk* to this new block



The *realloc()* function allocates memory from the heap.



---

Because it's possible that a new block will be allocated, any pointers into the old memory could be invalidated. These pointers will point to freed memory, with possible disastrous results, when a new block is allocated.

---

The *realloc()* function returns NULL when the memory pointed to by *old\_blk* can't be reallocated. In this case, the memory pointed to by *old\_blk* isn't freed, so be careful to maintain a pointer to the old memory block so it can be freed later.

In the following example, *buffer* is set to NULL if the function fails, and won't point to the old memory block. If *buffer* is your only pointer to the memory block, then you have lost access to this memory.

```
buffer = (char*)realloc(buffer, 100);
```

## Returns:

A pointer to the start of the allocated memory, or NULL if an error occurred (*errno* is set).

## Errors:

ENOMEM      Not enough memory.

EOK          No error.

## Examples:

```
#include <stdlib.h>
#include <malloc.h>

int main(void)
{
 char* buffer;
 char* new_buffer;

 buffer = (char*)malloc(80);
```

```
new_buffer = (char*)realloc(buffer, 100);
if(new_buffer == NULL) {
 /* not able to allocate larger buffer */

 return EXIT_FAILURE;
} else {
 buffer = new_buffer;
}

return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*calloc()*, *free()*, *malloc()*, *sbrk()*

## Synopsis:

```
#include <stdlib.h>

char * realpath(const char * pathname,
 char * resolved_name);
```

## Arguments:

|                      |                                                                       |
|----------------------|-----------------------------------------------------------------------|
| <i>pathname</i>      | The path name that you want to resolve.                               |
| <i>resolved_name</i> | A pointer to a buffer where the function can store the resolved name. |

## Library:

libc

## Description:

The *realpath()* function resolves all symbolic links, extra slash (/) characters and references to ./ and ../ in *pathname*, and copies the resulting absolute pathname into the memory referenced by *resolved\_name*.

To determine the size of the buffer pointed to by *resolved\_name*, call *fpathconf()* or *pathconf()* with an argument of `_PC_PATH_MAX`.

This function resolves both absolute and relative paths and returns the absolute pathname corresponding to *pathname*. All but the last component of *pathname* must exist when *realpath()* is called.

## Returns:

A pointer to *resolved\_name*, or NULL if an error occurs, in which case *resolved\_name* contains the pathname that caused the problem.

**Errors:**

The *realpath()* function may fail and set the external variable *errno* for any of the errors specified for the library functions *chdir()*, *close()*, *lstat()*, *open()*, *readlink()* and *getcwd()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

This implementation of *realpath()* differs slightly from the Solaris implementation. QNX always returns absolute pathnames, whereas the Solaris implementation, under certain circumstances, returns a relative *resolved\_name* when given a relative pathname.

**See also:**

*getcwd()*

**Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t recv(int s,
 void * buf,
 size_t len,
 int flags);
```

**Arguments:**

- |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>     | The descriptor for the socket; see <i>socket()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <i>buf</i>   | A pointer to a buffer where the function can store the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <i>len</i>   | The size of the buffer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>flags</i> | A combination formed by ORing one or more of the values: <ul style="list-style-type: none"><li>• MSG_OOB — process out-of-band data. This flag requests receipt of out-of-band data that wouldn't be received in the normal data stream. You can't use this flag with protocols that place expedited data at the head of the normal data queue.</li><li>• MSG_PEEK — peek at the incoming message. This flag causes the receive operation to return data from the beginning of the receive queue without removing that data from the queue. Thus, a subsequent receive call will return the same data.</li><li>• MSG_WAITALL — wait for full request or error. This flag requests that the operation block until the full request is satisfied. But the call may still return less data than requested if a signal is caught, if an error or disconnect occurs, or if the next data to be received is of a different type than that returned.</li></ul> |

---

☞ The MSG\_WAITALL flag isn't supported by the tiny TCP/IP stack. For more information, see `npm-ttcpip.so` in the *Utilities Reference*.

---

## Library:

`libsocket`

## Description:

The `recv()` function receives a message from a socket. It's normally used only on a *connected* socket — see `connect()` — and is identical to `recvfrom()` with a zero *from* parameter.

This routine returns the length of the message on successful completion. If a message is too long for the supplied buffer, *buf*, then excess bytes might be discarded, depending on the type of socket that the message is received from; see `socket()`.

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking — see `ioctl()` — in which case -1 is returned and the external variable *errno* is set to EWOULDBLOCK. Normally, the receive calls return any data available, up to the requested amount, rather than wait for the full amount requested; this behavior is affected by the socket-level options SO\_RCVLOWAT and SO\_RCVTIMEO described in `getsockopt()`.

You can use `select()` to determine when more data is to arrive.

## Returns:

The number of bytes received, or -1 if an error occurs (*errno* is set).

## Errors:

|        |                                                                                    |
|--------|------------------------------------------------------------------------------------|
| EBADF  | Invalid descriptor <i>s</i> .                                                      |
| EFAULT | The receive buffer is outside the process's address space.                         |
| EINTR  | The receive was interrupted by delivery of a signal before any data was available. |

ENOTCONN      The socket is associated with a connection-oriented protocol and hasn't been connected; see *connect()* and *accept()*.

EWOULDBLOCK  
Either the socket is marked nonblocking and the receive operation would block, or a receive timeout had been set and the timeout expired before data was received.

### Classification:

Standard Unix, POSIX 1003.1-2001

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*connect()*, *ioctl()*, *getsockopt()*, *read()*, *recvfrom()*, *recvmsg()*, *select()*, *socket()*

## ***recvfrom()***

© 2004, QNX Software Systems Ltd.

*Receive a message from the socket at a specified address*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t recvfrom(int s,
 void * buff,
 size_t len,
 int flags,
 struct sockaddr * from,
 socklen_t * fromlen);
```

### **Arguments:**

- |              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>     | The descriptor for the socket; see <i>socket()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <i>buf</i>   | A pointer to a buffer where the function can store the message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <i>len</i>   | The size of the buffer.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <i>flags</i> | A combination formed by ORing one or more of the values: <ul style="list-style-type: none"><li>• MSG_OOB — process out-of-band data. This flag requests receipt of out-of-band data that wouldn't be received in the normal data stream. You can't use this flag with protocols that place expedited data at the head of the normal data queue.</li><li>• MSG_PEEK — peek at the incoming message. This flag causes the receive operation to return data from the beginning of the receive queue without removing that data from the queue. Thus, a subsequent receive call will return the same data.</li><li>• MSG_WAITALL — wait for full request or error. This flag requests that the operation block until the full request is satisfied. But the call may still return less data than requested if a signal is caught, if an error or</li></ul> |



disconnect occurs, or if the next data to be received is of a different type than that returned.

---

☞ The MSG\_WAITALL flag isn't supported by the tiny TCP/IP stack. For more information, see `npm-ttcpip.so` in the *Utilities Reference*.

---

*from* NULL, or a pointer to a `sockaddr` object where the function can store the source address of the message.

*fromlen* A pointer to a `socklen_t` object that specifies the size of the *from* buffer. The function stores the actual size of the address in this object.

## Library:

`libsocket`

## Description:

The *recvfrom()* routine receives a message from the socket, *s*, whether or not it's connection-oriented.

If *from* is nonzero, and the socket is connectionless, the source address of the message is filled in. The parameter *fromlen* is a value-result parameter, initialized to the size of the buffer associated with *from*, and modified on return to indicate the actual size of the stored address.

This routine returns the length of the message on successful completion. If a message is too long for the supplied buffer, *buf*, excess bytes may be discarded depending on the type of socket that the message is received from — see *socket()*.

If no messages are available at the socket, the receive call waits for a message to arrive, unless the socket is nonblocking — see *ioctl()* — in which case *recvfrom()* returns -1 is returned and sets the external variable *errno* to EWOULDBLOCK. Normally, the receive calls return any data available, up to the requested amount, rather than wait for the full amount requested; this behavior is affected by the socket-level options SO\_RCVLOWAT and SO\_RCVTIMEO described in *getsockopt()*.

You can use *select()* to determine when more data is to arrive.

**Returns:**

The number of bytes received, or -1 if an error occurs (*errno* is set).

**Errors:**

- EBADF            Invalid descriptor *s*.
- EFAULT          The receive buffer pointer(s) point outside the process's address space.
- EINTR            The receive was interrupted by delivery of a signal before any data was available.
- ENOTCONN        The socket is associated with a connection-oriented protocol and hasn't been connected; see *connect()* and *accept()*.
- EWouldBlock     Either the socket is marked nonblocking and the receive operation would block, or a receive timeout had been set and the timeout expired before data was received.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*recv(), recvmsg(), select()*

## ***recvmsg()***

© 2004, QNX Software Systems Ltd.

*Receive a message and its header from a socket*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t recvmsg(int s,
 struct msghdr * msg,
 int flags);
```

### **Arguments:**

- s*            The descriptor for the socket; see *socket()*.
- msg*          A pointer to a **msghdr** structure where the function can store the message header; see below.
- len*          The size of the buffer.
- flags*        A combination formed by ORing one or more of the values:
- MSG\_OOB — process out-of-band data. This flag requests receipt of out-of-band data that wouldn't be received in the normal data stream. You can't use this flag with protocols that place expedited data at the head of the normal data queue.
  - MSG\_PEEK — peek at the incoming message. This flag causes the receive operation to return data from the beginning of the receive queue without removing that data from the queue. Thus, a subsequent receive call will return the same data.
  - MSG\_WAITALL — wait for full request or error. This flag requests that the operation block until the full request is satisfied. But the call may still return less data than requested if a signal is caught, if an error or disconnect occurs, or if the next data to be received is of a different type than that returned.

- 
- ☞ The MSG\_WAITALL flag isn't supported by the tiny TCP/IP stack. For more information, see `npm-ttcpip.so` in the *Utilities Reference*.
- 

## Library:

`libsocket`

## Description:

The *recvmsg()* routine receives a message from a socket, *s*, whether or not it's connection-oriented.

The *recvmsg()* call uses a `msg_hdr` structure to minimize the number of directly supplied parameters. This structure, defined in `<sys/socket.h>`, has the following form:

```
struct msg_hdr {
 caddr_t msg_name; /* optional address */
 u_int msg_namelen; /* size of address */
 struct iovec *msg_iov; /* scatter/gather array */
 u_int msg_iovlen; /* # elements in msg_iov */
 caddr_t msg_control; /* ancillary data, see below */
 u_int msg_controllen; /* ancillary data buffer len */
 int msg_flags; /* flags on received message */
};
```

The *msg\_name* and *msg\_namelen* parameters specify the address (source address for *recvmsg()*; destination address for *sendmsg()*) if the socket is unconnected; the *msg\_name* parameter may be given as a null pointer if no names are desired or required.

The *msg\_iov* and *msg\_iovlen* parameters describe scatter-gather locations, as discussed in *read()*.

The *msg\_control* parameter, whose length is determined by *msg\_controllen*, points to a buffer for other protocol-control related messages or for other miscellaneous ancillary data. The messages are of the form:

```
struct cmsghdr {
 u_int cmsgh_len; /* data byte count, including hdr */
 int cmsgh_level; /* originating protocol */
};
```

```

 int cmsg_type; /* protocol-specific type */
 /* followed by u_char cmsg_data[]; */
 };

```



Currently, the tiny TCP/IP stack doesn't support ancillary data. The `msg_controllen` member of `struct msghdr` must be 0.

The `msg_flags` field is set on return according to the message received:

|            |                                                                                                                           |
|------------|---------------------------------------------------------------------------------------------------------------------------|
| MSG_CTRUNC | Indicates that some control data was discarded due to lack of space in the buffer for ancillary data.                     |
| MSG_EOR    | Indicates end-of-record; the data returned completed a record.                                                            |
| MSG_OOB    | Indicates that expedited or out-of-band data was received.                                                                |
| MSG_TRUNC  | Indicates that the trailing portion of a datagram was discarded because the datagram was larger than the buffer supplied. |

### Returns:

The number of bytes received, or -1 if an error occurs (`errno` is set).

### Errors:

ENOMEM Not enough memory.

### Classification:

Standard Unix, POSIX 1003.1-2001

#### Safety

Cancellation point Yes

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*recv()*, *recvfrom()*, *sendmsg()*

## ***regcomp()***

© 2004, QNX Software Systems Ltd.

*Compile a regular expression*

### **Synopsis:**

```
#include <regex.h>

int regcomp(regex_t * preg,
 const char * pattern,
 int cflags);
```

### **Arguments:**

- |                |                                                                                                                                                                                                                                                                                                                                                                    |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>preg</i>    | A pointer to a <b>regex_t</b> object where the function can store the compiled regular expression.                                                                                                                                                                                                                                                                 |
| <i>pattern</i> | The regular expression that you want to compile; see below.                                                                                                                                                                                                                                                                                                        |
| <i>cflags</i>  | A bitwise inclusive OR of zero or more of the following flags: <ul style="list-style-type: none"><li>• REG_EXTENDED — use Extended Regular Expressions.</li><li>• REG_ICASE — ignore differences in case.</li><li>• REG_NEWLINE — treat &lt;newline&gt; as a regular character.</li><li>• REG_NOSUB — report only success/failure in <i>regexexec()</i>.</li></ul> |

### **Library:**

libc

### **Description:**

The *regcomp()* function prepares the regular expression, *preg*, for use by the function *regexexec()*, from the specification *pattern* and *cflags*. The member *re\_nsub* of *preg* is set to the number of subexpressions in *pattern*.

The functions that deal with regular expressions (*regcomp()*, *regerror()*, *regexexec()*, and *regfree()*) support two classes of regular



expressions, the *Basic* and *Extended Regular* Expressions. These classes are rigorously defined in IEEE P1003.2, *Regular Expression Notation*.

## Basic Regular Expressions

The Basic Regular Expressions are composed of these terms:

|                    |                                                                               |
|--------------------|-------------------------------------------------------------------------------|
| <code>x\$</code>   | <code>x</code> at end of line ( <code>\$</code> must be the last term).       |
| <code>^x</code>    | <code>x</code> at beginning of line ( <code>^</code> must be first the term). |
| <code>x*</code>    | Zero or more occurrences of <code>x</code> .                                  |
| <code>.</code>     | Any single character (except newline).                                        |
| <code>c</code>     | The character <code>c</code> .                                                |
| <code>xc</code>    | <code>x</code> followed by the character <code>c</code> .                     |
| <code>cx</code>    | Character <code>c</code> followed by <code>x</code> .                         |
| <code>[cd]</code>  | The characters <code>c</code> or <code>d</code> .                             |
| <code>[c-d]</code> | All characters between <code>c</code> and <code>d</code> , inclusive.         |
| <code>[^c]</code>  | Any character but <code>c</code> .                                            |

`[:classname:]`

Any of the following classes:

- `alnum`
- `alpha`
- `cntrl`
- `digit`
- `graph`
- `lower`
- `print`
- `punct`

- **space**
- **upper**
- **xdigit**

|                      |                                                        |
|----------------------|--------------------------------------------------------|
| <code>[[=c=]]</code> | All character in the equivalence class with <i>c</i> . |
| <code>[[.=.]]</code> | All collating elements.                                |
| <code>x{m,n}</code>  | <i>m</i> through <i>n</i> occurrences of <i>x</i> .    |
| <code>\c</code>      | Character <i>c</i> , even if <i>c</i> is an operator.  |
| <code>\(x\)</code>   | A labeled subexpression, <i>x</i> .                    |
| <code>\m</code>      | The <i>m</i> th subexpression encountered.             |
| <code>xy</code>      | Expression <i>x</i> followed by <i>y</i> .             |

### Extended Regular Expressions

The Extended Regular Expressions also include:

|                  |                                                   |
|------------------|---------------------------------------------------|
| <code>x+</code>  | One or more occurrences of <i>x</i> .             |
| <code>x?</code>  | Zero or one occurrences of <i>x</i> .             |
| <code>(x)</code> | Subexpression <i>x</i> (for precedence handling). |
| <code>x/y</code> | Expression <i>x</i> OR <i>y</i> .                 |

### Returns:

|     |                                                                  |
|-----|------------------------------------------------------------------|
| 0   | Success.                                                         |
| <>0 | An error occurred (use <i>regerror()</i> to get an explanation). |

**Examples:**

```

/*
 The following example prints out all lines
 from FILE "f" that match "pattern".
*/
#include <stdio.h>
#include <regex.h>
#include <limits.h>

#define BUFFER_SIZE 512

void grep(char* pattern, FILE* f)
{
 int t;
 regex_t re;
 char buffer[BUFFER_SIZE];

 if ((t=regcomp(&re, pattern, REG_NOSUB)) != 0) {
 regerror(t, &re, buffer, sizeof buffer);
 fprintf(stderr,"grep: %s (%s)\n",buffer,pattern);
 return;
 }
 while(fgets(buffer, BUFFER_SIZE, f) != NULL) {
 if(regexec(&re, buffer, 0, NULL, 0) == 0) {
 fputs(buffer, stdout);
 }
 }
 regfree(&re);
}

```

**Classification:**

POSIX 1003.1a

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Contributing author:**

Henry Spencer. For copyright information, see Third-Party Copyright Notices in this reference.

**See also:**

*regerror()*, *regexexec()*, *regfree()*

## Synopsis:

```
#include <regex.h>

size_t regerror(int err,
 const regex_t * reg,
 char * buf,
 size_t len);
```

## Arguments:

- err* The value returned by a previous call to *regcomp()* or *regexexec()*.
- reg* A pointer to the **regex\_t** object for the regular expression that you provided to the failed call to *regcomp()* or *regexexec()*.
- buf* A pointer to a buffer where the function can store the explanation.
- len* The length of the buffer, in characters.

## Library:

**libc**

## Description:

The *regerror()* function provides a string explaining an error code returned by *regcomp()* or *regexexec()*. The string is copied into *buf* for up to *len* characters.

## Returns:

The number of characters copied into the buffer.

**Examples:**

See *regcomp()*.

**Classification:**

POSIX 1003.1a

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Contributing author:**

Henry Spencer. For copyright information, see Third-Party Copyright Notices in this reference.

**See also:**

*regcomp()*, *regexec()*, *regfree()*

**Synopsis:**

```
#include <regex.h>

int regexec(const regex_t * preg,
 const char * string,
 size_t nmatch,
 regmatch_t * pmatch,
 int eflags);
```

**Arguments:**

- preg* A pointer to the **regex\_t** object for the regular expression that you want to execute. You must have compiled the expression by calling *regcomp()*.
- string* The string that you want to match against the regular expression.
- nmatch* The maximum number of matches to record in *pmatch*.
- pmatch* An array of **regmatch\_t** objects where the function can record the matches; see below.
- eflags* Execution parameters to *regexec()*. For example, you may need to call *regexec()* multiple times if the line you're processing is too large to fit into *string*. The *eflags* argument is the bitwise inclusive OR of zero or more of the following flags:
- REG\_NOTBOL — the *string* argument doesn't point to the beginning of a line.
  - REG\_NOTEOL — the end of *string* isn't the end of a line.

**Library:**

**libc**

## Description:

The *regexec()* function compares *string* against the compiled regular expression *preg*. If *regexec()* finds a match it returns zero; otherwise, it returns nonzero.

The *preg* argument represents a compiled form of either a Basic Regular Expression or Extended Regular Expression. These classes are rigorously defined in IEEE P1003.2, *Regular Expression Notation*, and are summarized in the documentation for *regcomp()*.

The *regexec()* function records the matches in the *pmatch* array, with *nmatch* specifying the maximum number of matches to record. The **regmatch\_t** structure contains the pointer members *rm\_sp* and *rm\_ep*, which are updated to identify the start and end of each matched substring. The pointers in *pmatch[0]* identify the substring corresponding to the entire expression, while those in *pmatch[1...nmatch]* identify up to the first *nmatch* subexpressions. Unused elements of the *pmatch* array are set to NULL.



---

You can disable the recording of substrings by either specifying REG\_NOSUB in *regcomp()*, or by setting *nmatch* to zero.

---

## Returns:

- 0        The *string* argument matches *preg*.
- <>0     A match wasn't found, or an error occurred (use *regerror()* to get an explanation).

## Examples:

See *regcomp()*.

## Classification:

POSIX 1003.1a



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Contributing author:**

Henry Spencer. For copyright information, see Third-Party Copyright Notices in this reference.

**See also:**

*regcomp()*, *regerror()*, *regfree()*

## ***regfree()***

© 2004, QNX Software Systems Ltd.

*Release memory allocated for a regular expression*

### **Synopsis:**

```
#include <regex.h>

void regfree(regex_t * preg);
```

### **Arguments:**

*preg* A pointer to the **regex\_t** object for the regular expression that you want to free; see *regcomp()*.

### **Library:**

libc

### **Description:**

The *regfree()* function releases all memory allocated by *regcomp()* associated with *preg*.

### **Examples:**

See *regcomp()*.

### **Classification:**

POSIX 1003.1a

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Contributing author:**

Henry Spencer. For copyright information, see Third-Party Copyright Notices in this reference.

**See also:**

*regcomp()*, *regerror()*, *regexec()*

## ***remainder()*, *remainderf()***

© 2004, QNX Software Systems Ltd.

*Compute the floating point remainder*

### **Synopsis:**

```
#include <math.h>

double remainder(double x,
 double y);

float remainderf(float x,
 float y);
```

### **Arguments:**

*x*     The numerator of the division.  
*y*     The denominator.

### **Library:**

`libm`

### **Description:**

The *remainder()* and *remainderf()* functions return the floating point remainder  $r = x - ny$ , where  $y$  is nonzero. The value  $n$  is the integral value nearest the exact value  $x/y$ . When  $|n - x/y| = \frac{1}{2}$ , the value  $n$  is chosen to be even.

The behavior of *remainder()* is independent of the rounding mode.

### **Returns:**

The floating point remainder  $r = x - ny$ , where  $y$  is nonzero.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:***drem()*, *modf()*

## ***remove()***

© 2004, QNX Software Systems Ltd.

*Remove a link to a file*

### **Synopsis:**

```
#include <stdio.h>

int remove(const char * filename);
```

### **Arguments:**

*filename*     The path to the file that you want to delete.

### **Library:**

`libc`

### **Description:**

The *remove()* function removes a link to a file:

- If the *filename* names a symbolic link, *remove()* removes the link, but doesn't affect the file or directory that the link goes to.
- If the *filename* isn't a symbolic link, *remove()* removes the link and decrements the link count of the file that the link refers to.

If the link count of the file becomes zero, and no process has the file open, then the space that the file occupies is freed, and no one can access the file anymore.

If one or more processes have the file open when the last link is removed, the link is removed, but the removal of the file is delayed until all references to it have been closed.

This function is equivalent to *unlink()*.



---

To remove a directory, call *rmdir()*.

---

**Returns:**

- |         |                                              |
|---------|----------------------------------------------|
| 0       | The operation was successful.                |
| Nonzero | The operation failed ( <i>errno</i> is set). |

**Errors:**

- |              |                                                                                                                                                                                                                        |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of <i>filename</i> , or write permission is denied on the directory containing the link to be removed.                                                                     |
| EBUSY        | The directory named by <i>filename</i> can't be unlinked because it's being used by the system or another process, and the target filesystem or resource manager considers this to be an error.                        |
| ENAMETOOLONG | The <i>filename</i> argument exceeds <code>PATH_MAX</code> in length, or a pathname component is longer than <code>NAME_MAX</code> .                                                                                   |
| ENOENT       | The named file doesn't exist, or <i>filename</i> is an empty string.                                                                                                                                                   |
| ENOSYS       | The <i>remove()</i> function isn't implemented for the filesystem specified by <i>filename</i> .                                                                                                                       |
| ENOTDIR      | A component of <i>filename</i> isn't a directory.                                                                                                                                                                      |
| EPERM        | The file named by <i>filename</i> is a directory, and either the calling process doesn't have the appropriate privileges, or the target filesystem or resource manager prohibits using <i>remove()</i> on directories. |
| EROFS        | The directory entry to be unlinked resides on a read-only filesystem.                                                                                                                                                  |

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 if(remove("vm.tmp")) {
 puts("Error removing vm.tmp!");
 return EXIT_FAILURE;
 }

 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, rmdir(), unlink()*



## Synopsis:

```
#include <stdio.h>

int rename(const char* old,
 const char* new);
```

## Arguments:

*old*      The path to the file that you want to rename.

*new*      The new name for the file.

## Library:

libc

## Description:

The *rename()* function renames the file indicated *old* to the name specified in *new*.

If a file (or empty directory) named *new* exists, it's overwritten.

## Returns:

0            Success.

Nonzero     An error occurred (*errno* is set ).

## Errors:

EACCESS     The calling program doesn't have permission to search one of the components of either path prefix, or one of the directories containing *old* or *new* denies write permission.

EBUSY       The directory named by *old* or *new* can't be renamed because it is in use by another process.

|              |                                                                                                                              |
|--------------|------------------------------------------------------------------------------------------------------------------------------|
| EEXIST       | The file specified by <i>new</i> is directory that contains files.                                                           |
| EINVAL       | The <i>new</i> directory pathname contains the <i>old</i> directory.                                                         |
| EISDIR       | The file specified by <i>new</i> is a directory and <i>old</i> is a file.                                                    |
| ELOOP        | Too many levels of symbolic links.                                                                                           |
| EMLINK       | The file named by <i>old</i> is a directory, and the link count of the parent directory of <i>new</i> would exceed LINK_MAX. |
| ENAMETOOLONG | The length of <i>old</i> or <i>new</i> exceeds PATH_MAX.                                                                     |
| ENOENT       | The <i>old</i> file doesn't exist, or <i>old</i> or <i>new</i> is an empty string.                                           |
| ENOSPC       | The directory that would contain <i>new</i> can't be extended.                                                               |
| ENOSYS       | The <i>rename()</i> function isn't implemented for the filesystem specified in <i>old</i> or <i>new</i> .                    |
| ENOTDIR      | A component of either path prefix isn't a directory, or <i>old</i> is a directory and <i>new</i> isn't.                      |
| ENOTEMPTY    | The file specified by <i>new</i> is a directory that contains files.                                                         |
| EROFS        | The <i>rename()</i> would affect files on a read-only filesystem.                                                            |
| EXDEV        | The files or directories named by <i>old</i> and <i>new</i> are on different filesystems.                                    |

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 if(rename("old.dat", "new.dat")) {
 puts("Error renaming old.dat to new.dat.");

 return EXIT_FAILURE;
 }

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*

## ***res\_init()***

© 2004, QNX Software Systems Ltd.

*Initialize the Internet domain name resolver routines*

### **Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_init(void);
```

### **Library:**

libsocket

### **Description:**

The resolver routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers.

The *res\_init()* routine reads the resolver configuration file (if one is present; see */etc/resolv.conf* in the *Utilities Reference*) to get the default domain name, search list, and Internet address of the local name servers. If no server is configured, the host running the resolver is tried. If not specified in the configuration file, the current domain name is defined by the hostname; the domain name can be overridden by the environment variable **LOCALDOMAIN**. Initialization normally occurs on the first call to one of the resolver routines.

### **Resolver configuration**

Global configuration and state information used by these routines is kept in the `__res_state` structure `_res`, which is defined in `<resolv.h>`. Since most of the values have reasonable defaults, you can generally ignore them.

The `_res.options` member is a simple bit mask that contains the bitwise OR of the enabled options. The following options are defined in `<resolv.h>`:

RES\_DEBUG      Print debugging messages.

**RES\_DEFNAMES**

If this option is set, *res\_search()* appends the default domain name to single-component names (those that don't contain a dot). This option is enabled by default.

**RES\_DNSRCH**

If this option is set, *res\_search()* searches for hostnames in the current domain and in parent domains. This is used by the standard host lookup routine, *gethostbyname()*. This option is enabled by default.

**RES\_INIT**

True if the initial name server address and default domain name are initialized (i.e. *res\_init()* has been called).

**RES\_RECURSE**

Set the recursion-desired bit in queries. This is the default. Note that *res\_send()* doesn't do iterative queries — it expects the name server to handle recursion.

**RES\_STAYOPEN**

Used with **RES\_USEVC** to keep the TCP connection open between queries. This is useful only in programs that regularly do many queries. UDP should be the mode you normally use.

**RES\_USEVC**

Instead of UDP datagrams, use TCP connections for queries.

**Returns:**

- |         |                    |
|---------|--------------------|
| 0       | Success.           |
| Nonzero | An error occurred. |

**Errors:**

See *herror()*.

**Files:**

`/etc/resolv.conf`  
Resolver configuration file.

**Environment variables:**

**LOCALDOMAIN**

When set, **LOCALDOMAIN** contains a domain name that overrides the current domain name.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*dn\_comp()*, *dn\_expand()*, *gethostbyname()*, *res\_mkquery()*,  
*res\_query()*, *res\_querydomain()*, *res\_search()*, *res\_send()*

**hostname**, `/etc/resolv.conf` in the *Utilities Reference*

*RFC 974*, *RFC 1032*, *RFC 1033*, *RFC 1034*, *RFC 1035*

**Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_mkquery(int op,
 const char * dname,
 int class,
 int type,
 const u_char * data,
 int datalen,
 const u_char * newrr,
 u_char * buf,
 int buflen);
```

**Arguments:**

- op* Usually QUERY, but it can be also IQUERY or NS\_NOTIFY\_OP. Note that not all of the query types defined in <arpa/nameser.h> are supported.
- dname* The domain name for the query.
- class* The class of information that you want; one of:
- C\_IN — ARPA Internet.
  - C\_CHAOS — Chaos net (MIT).
  - C\_HS — Hesiod name server (MIT).
  - C\_ANY — any class.
- You typically use C\_IN.
- type* The type of information that you want. You typically use T\_PTR, but you can use any of the T\_\* constants defined in <arpa/nameser.h>.
- data* NULL, or a pointer to resource record data.
- datalen* The length of the data.

*newrr* Currently unused. This argument is intended for making update messages.

*buf* A pointer to a buffer where the function can build the query.

*buflen* The length of the buffer.

## Library:

`libsocket`

## Description:

The *res\_mkquery()* function is a low-level routine that's used by *res\_query()* to construct an Internet domain name query. This routine constructs a standard query message and places it in *buf*. It returns the size of the query, or -1 if the query is larger than *buflen*.

The resolver routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers. Global configuration and state information used by the resolver routines is kept in the structure `_res`. For more information on the options, see *res\_init()*.

## Returns:

The size of the prepared query, in bytes, or -1 if an error occurs.

## Files:

`/etc/resolv.conf`  
Resolver configuration file.

## Environment variables:

### **LOCALDOMAIN**

When set, **LOCALDOMAIN** contains a domain name that overrides the current domain name.



## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*dn\_comp(), dn\_expand(), gethostbyname(), res\_init(), res\_query(),  
res\_querydomain(), res\_search(), res\_send()*

*hostname, /etc/resolv.conf* in the *Utilities Reference*

*RFC 974, RFC 1032, RFC 1033, RFC 1034, RFC 1035*

## ***res\_query()***

© 2004, QNX Software Systems Ltd.

*Query the local Internet domain name server*

### **Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_query(const char * dname,
 int class,
 int type,
 u_char * answer,
 int anslen);
```

### **Arguments:**

*dname*      The fully qualified domain name that you want to query.

*class*      The class of information that you want; one of:

- C\_IN — ARPA Internet.
- C\_CHAOS — Chaos net (MIT).
- C\_HS — Hesiod name server (MIT).
- C\_ANY — any class.

You typically use C\_IN.

*type*      The type of information that you want. You typically use T\_PTR, but you can use any of the T\_\* constants defined in `<arpa/nameser.h>`.

*answer*    A pointer to a buffer where the function can store the answer to the query.

*anslen*    The length of the buffer.

### **Library:**

`libsocket`

## Description:

The *res\_query()* function provides an interface to the server query mechanism. It constructs a query, sends it to the local server, waits for a response, and makes preliminary checks on the reply. The query requests information of the specified *type* and *class* for the specified fully qualified domain name *dname*. The reply message is left in the *answer* buffer with length *anslen* supplied by the caller.

The resolver routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers. Global configuration and state information used by the resolver routines is kept in the structure *\_res*. For more information on the options, see *res\_init()*.

The *res\_query()* function uses the following lower-level routines:

- *res\_mkquery()* constructs a standard query message.
- *res\_send()* sends the preformatted query and returns an answer.
- *dn\_comp()* compresses a domain name.
- *dn\_expand()* expands the compressed domain name to a full domain name.

## Returns:

The length of a reply message, in bytes, or -1 if an error occurs (*h\_errno* is set).

## Errors:

See *herror()*.

## Files:

`/etc/resolv.conf`

Resolver configuration file.

**Environment variables:**

**LOCALDOMAIN**

When set, **LOCALDOMAIN** contains a domain name that overrides the current domain name.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*dn\_comp(), dn\_expand(), gethostbyname(), perror(), res\_init(), res\_mkquery(), res\_querydomain(), res\_search(), res\_send()*

*hostname, /etc/resolv.conf* in the *Utilities Reference*

*RFC 974, RFC 1032, RFC 1033, RFC 1034, RFC 1035*

### **Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_querydomain(const char * name,
 const char * domain,
 int class,
 int type,
 u_char * answer,
 int anslen);
```

### **Arguments:**

*name*      The host name that you want to query.

*domain*    The domain name that you want to query.

*class*     The class of information that you want; one of:

- C\_IN — ARPA Internet.
- C\_CHAOS — Chaos net (MIT).
- C\_HS — Hesiod name server (MIT).
- C\_ANY — any class.

You typically use C\_IN.

*type*      The type of information that you want. You typically use T\_PTR, but you can use any of the T\_\* constants defined in <arpa/nameser.h>.

*answer*    A pointer to a buffer where the function can store the answer to the query.

*anslen*    The length of the buffer.

## Library:

`libsocket`

## Description:

The *res\_querydomain()* function provides an interface to the server query mechanism. It constructs a query, sends it to the local server, waits for a response, and makes preliminary checks on the reply. The query requests information of the specified *type* and *class* for the host specified by concatenating *name* and *domain*. The trailing dot is removed from *name* if *domain* is 0.

The reply message is left in the *answer* buffer with length *anslen* supplied by the caller.

## Returns:

0      Success.  
-1     An error occurred.

## Files:

`/etc/resolv.conf`  
Resolver configuration file.

## Environment variables:

**LOCALDOMAIN**

When set, **LOCALDOMAIN** contains a domain name that overrides the current domain name.

## Classification:

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*res\_init()*, *res\_query()*

## ***res\_search()***

© 2004, QNX Software Systems Ltd.

*Query a local server, using search options*

### **Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_search(const char * dname,
 int class,
 int type,
 u_char * answer,
 int anslen);
```

### **Arguments:**

*dname*      The fully qualified domain name that you want to query.

*class*      The class of information that you want; one of:

- C\_IN — ARPA Internet.
- C\_CHAOS — Chaos net (MIT).
- C\_HS — Hesiod name server (MIT).
- C\_ANY — any class.

You typically use C\_IN.

*type*        The type of information that you want. You typically use T\_PTR, but you can use any of the T\_\* constants defined in <arpa/nameser.h>.

*answer*     A pointer to a buffer where the function can store the answer to the query.

*anslen*     The length of the buffer.

### **Library:**

libsocket



## Description:

The *res\_search()* routine makes an Internet domain name search. Like *res\_query()*, *res\_search()* makes a query and waits for a response. But it also implements the default and search rules controlled by the `RES_DEFNAMES` and `RES_DNSRCH` options. It returns the first successful reply.

The resolver routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers.

Global configuration and state information used by the resolver routines is kept in the structure `_res`. For more information on the options, see *res\_init()*.

## Returns:

The length of a reply message, in bytes, or -1 if an error occurs (`h_errno` is set).

## Errors:

See *herror*.

## Files:

`/etc/resolv.conf`

Resolver configuration file.

## Environment variables:

`LOCALDOMAIN`

When set, `LOCALDOMAIN` contains a domain name that overrides the current domain name.

## Classification:

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*dn\_comp(), dn\_expand(), gethostbyname(), herror(), res\_init(),  
res\_mkquery(), res\_query(), res\_querydomain(), res\_send()*

**hostname**, **/etc/resolv.conf** in the *Utilities Reference*

*RFC 974, RFC 1032, RFC 1033, RFC 1034, RFC 1035*

**Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <arpa/nameser.h>
#include <resolv.h>

int res_send(const u_char * msg,
 int msglen,
 u_char * answer,
 int anslen);
```

**Arguments:**

|               |                                                                             |
|---------------|-----------------------------------------------------------------------------|
| <i>msg</i>    | The preformatted Internet domain name query that you want to send.          |
| <i>msglen</i> | The length of the message.                                                  |
| <i>answer</i> | A pointer to a buffer where the function can store the answer to the query. |
| <i>anslen</i> | The length of the buffer.                                                   |

**Library:**

**libsocket**

**Description:**

The *res\_send()* function is a low-level routine that's used by *res\_query()* to send a preformatted Internet domain name query and return an answer. It calls *res\_init()* if RES\_INIT isn't set, sends the query to the local name server, and handles timeouts and retries.

The resolver routines are used for making, sending, and interpreting query and reply messages with Internet domain name servers.

Global configuration and state information used by the resolver routines is kept in the structure *\_res*. For more information on the options, see *res\_init()*.

**Returns:**

The length of a reply message, in bytes; or -1 if an error occurs.

**Errors:**

ECONNREFUSED

No name servers found.

ETIMEDOUT

No answer obtained.

**Files:**

`/etc/resolv.conf`

Resolver configuration file.

**Environment variables:**

LOCALDOMAIN

When set, **LOCALDOMAIN** contains a domain name that overrides the current domain name.

**Classification:**

Unix

**Safety**

---

Cancellation point    Yes

Interrupt handler    No

Signal handler    No

Thread    No

**See also:**

*dn\_comp(), dn\_expand(), gethostbyname(), res\_init(), res\_mkquery(),  
res\_query(), res\_querydomain(), res\_search()*

**hostname**, **/etc/resolv.conf** in the *Utilities Reference*

*RFC 974, RFC 1032, RFC 1033, RFC 1034, RFC 1035*

## ***resmgr\_attach()***

© 2004, QNX Software Systems Ltd.

*Attach a path to the pathname space*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int resmgr_attach (
 dispatch_t *dpp,
 resmgr_attr_t *attr,
 const char *path,
 enum _file_type file_type,
 unsigned flags,
 const resmgr_connect_funcs_t *connect_funcs,
 const resmgr_io_funcs_t *io_funcs,
 RESMGR_HANDLE_T *handle);
```

### **Arguments:**

|                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>dpp</i>       | A dispatch handle created by <i>dispatch_create()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <i>attr</i>      | A pointer to a <b>resmgr_attr_t</b> structure that defines attributes for the resource manager; see below.                                                                                                                                                                                                                                                                                                                                                                     |
| <i>path</i>      | NULL, or the path that you want to attach the resource manager to; see below.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <i>file_type</i> | The file type; one of the following (defined in <b>&lt;sys/ftype.h&gt;</b> ): <ul style="list-style-type: none"><li>• <b>_FTYPE_ANY</b> — the path name can be anything.</li><li>• <b>_FTYPE_LINK</b> — reserved for the Process Manager.</li><li>• <b>_FTYPE_MOUNT</b> — receive mount requests on the path (<i>path</i> must be NULL).</li><li>• <b>_FTYPE_MQUEUE</b> — reserved for a mqueue manager.</li><li>• <b>_FTYPE_PIPE</b> — reserved for a pipe manager.</li></ul> |

- `_FTYPE_SEM` — reserved for a semaphore manager.
- `_FTYPE_SHMEM` — reserved for a shared memory object.
- `_FTYPE_SOCKET` — reserved for a socket manager.
- `_FTYPE_SYMLINK` — reserved for the Process Manager.

*flags*

Flags that control the pathname resolution:

- `_RESMGR_FLAG_AFTER`
- `_RESMGR_FLAG_BEFORE`
- `_RESMGR_FLAG_OPAQUE`
- `_RESMGR_FLAG_DIR`
- `_RESMGR_FLAG_FTYPEONLY`
- `_RESMGR_FLAG_MASK`
- `_RESMGR_FLAG_SELF`

For more information, see “Flags,” below.

*connect\_funcs*

A pointer to the `resmgr_connect_funcs_t` structure that defines the POSIX-level connect functions.

*io\_funcs*

A pointer to the `resmgr_io_funcs_t` structure that defines the POSIX-level I/O functions.

*handle*

A pointer to an arbitrary structure that you want to associate with the pathname you’re attaching. For most resource managers, this is an `iofunc_attr_t` structure.

**Library:**

libc

**Description:**

The *resmgr\_attach()* function puts the *path* into the general pathname space and binds requests on this path to the dispatch handle *dpp*.

Most of the above file types are used for special services that have their own open function associated with them. For example, the mqueue manager specifies *file\_type* as `_FTYPE_MQUEUE` and *mq\_open()* requests a pathname match of the same type.

Specify `_FTYPE_ANY` for normal filesystems and simple devices, such as serial ports, that don't have their own special open type. Also if you can handle the type of service or a redirection node to a manager that does. Most resource managers are of this type.

Your resource manager won't receive messages from an open of an inappropriate type. The following table shows the different open function types and the types of pathnames they'll match.

| <b>Function:</b>  | <b><i>file_type:</i></b>   | <b>Matches pathname of type:</b>                         |
|-------------------|----------------------------|----------------------------------------------------------|
| <i>mq_open()</i>  | <code>_FTYPE_MQUEUE</code> | <code>_FTYPE_ANY</code><br><code>_FTYPE_MQUEUE</code>    |
| <i>open()</i>     | <code>_FTYPE_ANY</code>    | all types                                                |
| <i>pipe()</i>     | <code>_FTYPE_PIPE</code>   | <code>_FTYPE_ANY</code> or<br><code>_FTYPE_PIPE</code>   |
| <i>sem_open()</i> | <code>_FTYPE_SEM</code>    | <code>_FTYPE_ANY</code> or<br><code>_FTYPE_SEM</code>    |
| <i>shm_open()</i> | <code>_FTYPE_SHMEM</code>  | <code>_FTYPE_ANY</code> or<br><code>_FTYPE_SHMEM</code>  |
| <i>socket()</i>   | <code>_FTYPE_SOCKET</code> | <code>_FTYPE_ANY</code> or<br><code>_FTYPE_SOCKET</code> |



The generic *open()* can be used to open a pathname of any type.

If you want to use the POSIX functions, we've provided you with the POSIX layer; to fill your connect and I/O functions tables with the default handler functions supplied by the POSIX layer library, call *iofunc\_func\_init()*. You can then override the defaults placed in the structures with your own handlers.

In the most general case, the last argument, *handle* is an arbitrary structure that you wish to have associated with the pathname you're attaching. Practically, however, we recommend that it contain the POSIX layer's well defined attributes structure, *iofunc\_attr\_t*, because this lets you use the POSIX layer default library. You can extend the data that's contained in the attributes structure to contain any device-specific data that you may require. This is commonly done, and is described in the "Extending Data Control Structures (DCS)" section in the Writing a Resource Manager chapter of the *Programmer's Guide*.

In order to use the POSIX layer default library, the attributes structure must be bound into the Open Control Block, and you must use the POSIX layer's *iofunc\_ocb\_t* OCB. This is described in the documentation for *resmgr\_open\_bind()*, as well as in the above reference.

### **resmgr\_attr\_t structure**

You can specify attributes such as the maximum message size, number of parts (number of IOVs in context), and flags in the *attr* structure. The *resmgr\_attr\_t* structure looks like this:

```
typedef struct _resmgr_attr {
 unsigned flags;
 unsigned nparts_max;
 unsigned msg_max_size;
 int (*other_func)
 (resmgr_context_t *, void *msg);
} resmgr_attr_t;
```

The members include:

|                     |                                                                                                                                                                                                                                                             |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>nparts_max</i>   | The number of components to allocate for the IOV array. If you specify 0, the resource manager library bumps the value to the minimum usable by the library itself.                                                                                         |
| <i>msg_max_size</i> | The minimum amount of room to reserve for receiving a message that's allocated in <i>resmgr_context_alloc()</i> . If the value is too low, or you specify it as 0, <i>resmgr_attach()</i> picks a value that's usable.                                      |
| <i>other_func</i>   | A pointer to a function that's called if the resource manager receives an I/O message that it didn't successfully handle. This function is attached only if the RESMGR_FLAG_ATTACH_OTHERFUNC flag (defined in <code>&lt;sys/dispatch.h&gt;</code> ) is set. |

## Flags

The *flags* argument specifies additional information to control the pathname resolution. The flags (defined in `<sys/resmgr.h>`) include at least the following bits:

### `_RESMGR_FLAG_AFTER`

Force the path to be resolved after others with the same pathname at the same mountpoint.

### `_RESMGR_FLAG_BEFORE`

Force the path to be resolved before others with the same pathname at the same mountpoint.

### `_RESMGR_FLAG_DIR`

Treat the pathname as a directory and allow the resolving of longer pathnames. The `_IO_CONNECT` message contains the pathname passed to the client *open()* with the matching prefix stripped off. Without this flag, the pathname is treated as a simple file requiring an exact match.

| <b>Attached path</b> | <b>Opened path</b> | <b>_RESMGR_FLAG_DIR set</b> | <b>_RESMGR_FLAG_DIR clear</b> |
|----------------------|--------------------|-----------------------------|-------------------------------|
| /a/b                 | /a/b               | match ""                    | match ""                      |
| /a/b                 | /a/b/c             | match c                     | no match                      |
| /a/b                 | /a/b/c/d           | match c/d                   | no match                      |
| /a/b                 | /a/bc              | no match                    | no match                      |

You can't attach a directory pathname that contains, as a subset, an existing file pathname. Likewise, you can't attach a file pathname that's a subset of an existing directory pathname.

| <b>Existing path</b> | <b>New path</b> | <b>New path allowed?</b> |
|----------------------|-----------------|--------------------------|
| dir /a/b             | dir /a          | yes                      |
| dir /a/b             | dir /a/b/c      | yes                      |
| file /a/b            | dir /a          | yes                      |
| file /a/b            | dir /a/b/c      | no, dir beneath a file   |
| dir /a/b             | file /a         | no, dir beneath a file   |
| dir /a/b             | file /a/b/c     | yes                      |
| file /a/b            | file /a         | yes                      |
| file /a/b            | file /a/b/c     | yes                      |

**\_RESMGR\_FLAG\_FTYPEONLY**

Handle only requests for the specific filetype indicated. The pathname must be NULL.

**\_RESMGR\_FLAG\_OPAQUE**

Don't resolve paths to mountpoints on a path shorter than this (i.e. find the longest match against all pathnames attached).

**\_RESMGR\_FLAG\_SELF**

Allow requests to resolve back to this server (a deadlock is possible).

## Returns:

A unique link ID associated with this attach, or -1 on failure (*errno* is set).

The returned ID is needed to detach the pathname at a later time using *resmgr\_detach()*. The ID is also passed back in the *resmgr\_handler()* function in *ctp->id*.

## Errors:

- ENOMEM      There isn't enough free memory to complete the operation.
- ENOTDIR     A component of the pathname wasn't a directory entry.

## Examples:

Here's an example of a simple single-threaded resource manager:

```
#include <stdio.h>
#include <stddef.h>
#include <stdlib.h>
#include <sys/iofunc.h>
#include <sys/dispatch.h>

static resmgr_connect_funcs_t connect_funcs;
static resmgr_io_funcs_t io_funcs;
static iofunc_attr_t attr;

int main(int argc, char **argv)
{
 dispatch_t *dpp;
 resmgr_attr_t resmgr_attr;
 resmgr_context_t *ctp;
 int id;

 /* initialize dispatch interface */
 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 /* initialize resource manager attributes */
 memset(&resmgr_attr, 0, sizeof resmgr_attr);
 resmgr_attr.nparts_max = 1;
}
```

```

resmgr_attr.msg_max_size = 2048;

/* initialize functions for handling messages */
iofunc_func_init(_RESMGR_CONNECT_NFUNCS, &connect_funcs,
 _RESMGR_IO_NFUNCS, &io_funcs);

/* initialize attribute structure */
iofunc_attr_init(&attr, S_IFNAM | 0666, 0, 0);

/* attach our device name (passing in the POSIX defaults
 from the iofunc_func_init and iofunc_attr_init functions)
 */
if ((id = resmgr_attach
 (dpp, &resmgr_attr, "/dev/mynull", _FTYPE_ANY, 0,
 &connect_funcs, &io_funcs, &attr)) == -1) {
 fprintf(stderr, "%s: Unable to attach name.\n", \
 argv[0]);
 return EXIT_FAILURE;
}

/* allocate a context structure */
ctp = resmgr_context_alloc(dpp);

/* start the resource manager message loop */
while (1) {
 if ((ctp = resmgr_block(ctp)) == NULL) {
 fprintf(stderr, "block error\n");
 return EXIT_FAILURE;
 }
 resmgr_handler(ctp);
}
}

```

For more examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, and *thread\_pool\_create()*. For more information on writing a resource manager, see the “Writing a Resource Manager” chapter in the *Programmer’s Guide*.

## Classification:

QNX Neutrino

### Safety

Cancellation point Yes

*continued...*

## Safety

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

## **Caveats:**

If your application calls this function, it must run as **root**.

## **See also:**

*dispatch\_create()*, *iofunc\_attr\_init()*, **iofunc\_attr\_t**,  
*iofunc\_func\_init()*, **iofunc\_ocb\_t**, *resmgr\_block()*,  
**resmgr\_connect\_funcs\_t**, *resmgr\_context\_alloc()*,  
*resmgr\_context\_free()*, *resmgr\_detach()*, *resmgr\_handler()*,  
**resmgr\_io\_funcs\_t**

“Writing a Resource Manager” chapter of the *Programmer’s Guide*.

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

resmgr_context_t * resmgr_block
 (resmgr_context_t * ctp);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

### **Library:**

`libc`

### **Description:**

The `resmgr_block()` function waits for a message (created by a call to `resmgr_context_alloc()`) for context *ctp*.



---

This function is a special case of `dispatch_block()` that you should use only with a simple resource manager. If you need to attach pulses or other messages, then you should use `dispatch_block()`.

---

### **Returns:**

The same pointer as *ctp*, or NULL if an error occurs (*errno* is set).

### **Errors:**

EFAULT A fault occurred when the kernel tried to access the buffers provided. Because the OS accesses the sender's buffers only when `MsgReceive()` is called, a fault could occur *in the sender* if the sender's buffers are invalid. If a fault occurs when accessing

the sender buffers (only) they'll receive an EFAULT and the *MsgReceive()* won't unblock.

EINTR           The call was interrupted by a signal.

ETIMEDOUT       A kernel timeout (that was set with *dispatch\_timeout()*) unblocked the call.

## Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 dispatch_t *dpp;
 resmgr_context_t *ctp;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 :

 ctp = resmgr_context_alloc(dpp);

 while (1) {
 if ((ctp = resmgr_block(ctp)) == NULL) {
 fprintf(stderr, "block error\n");
 return EXIT_FAILURE;
 }
 resmgr_handler(ctp);
 }
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

Use *resmgr\_block()* only in a simple resource manager and when you don't use *message\_attach()*, *pulse\_attach()*, or *select\_attach()*.

**See also:**

*dispatch\_block()*, *resmgr\_attach()*, *resmgr\_context\_alloc()*,  
*resmgr\_handler()*

“Components of a Resource Manager” section of the Writing a Resource Manager chapter in the *Programmer's Guide*.

## **resmgr\_connect\_funcs\_t**

© 2004, QNX Software Systems Ltd.

*Table of POSIX-level connect functions*

### **Synopsis:**

```
#include <sys/resmgr.h>

typedef struct _resmgr_connect_funcs {

 unsigned nfuncs;

 int (*open) (resmgr_context_t *ctp, io_open_t *msg,
 RESMGR_HANDLE_T *handle, void *extra);

 int (*unlink) (resmgr_context_t *ctp, io_unlink_t *msg,
 RESMGR_HANDLE_T *handle, void *reserved);

 int (*rename) (resmgr_context_t *ctp, io_rename_t *msg,
 RESMGR_HANDLE_T *handle,
 io_rename_extra_t *extra);

 int (*mknod) (resmgr_context_t *ctp, io_mknod_t *msg,
 RESMGR_HANDLE_T *handle, void *reserved);

 int (*readlink) (resmgr_context_t *ctp, io_readlink_t *msg,
 RESMGR_HANDLE_T *handle, void *reserved);

 int (*link) (resmgr_context_t *ctp, io_link_t *msg,
 RESMGR_HANDLE_T *handle,
 io_link_extra_t *extra);

 int (*unblock) (resmgr_context_t *ctp, io_pulse_t *msg,
 RESMGR_HANDLE_T *handle, void *reserved);

 int (*mount) (resmgr_context_t *ctp, io_mount_t *msg,
 RESMGR_HANDLE_T *handle,
 io_mount_extra_t *extra);

} resmgr_connect_funcs_t;
```

### **Description:**

The `resmgr_connect_funcs_t` structure is a table of the POSIX-level connect functions that are used by a resource manager.

You can initialize this table by calling *iofunc\_func\_init()* and then overriding the defaults with your own functions.

This structure includes *nfuncs*, which indicates how many functions are in the table (in case the structure grows in the future), along with these functions:

| <b>Member:</b>  | <b>Used to:</b>                                 | <b>Default:</b>              |
|-----------------|-------------------------------------------------|------------------------------|
| <i>open</i>     | Handle <code>_IO.CONNECT</code> messages        | <i>iofunc_open_default()</i> |
| <i>unlink</i>   | Unlink the resource                             | None                         |
| <i>rename</i>   | Rename the resource                             | None                         |
| <i>mknod</i>    | Create a filesystem entry point                 | None                         |
| <i>readlink</i> | Read a symbolic link                            | None                         |
| <i>link</i>     | Create a symbolic link                          | None                         |
| <i>unblock</i>  | Unblock the resource if an operation is aborted | None                         |
| <i>mount</i>    | Mount a filesystem                              | None                         |

### **Classification:**

QNX Neutrino

### **See also:**

*iofunc\_func\_init()*, *iofunc\_open\_default()*, **resmgr\_io\_funcs\_t**

Writing a Resource Manager chapter of the QNX Neutrino  
*Programmer's Guide*

## ***resmgr\_context\_alloc()***

© 2004, QNX Software Systems Ltd.

*Allocate a resource-manager context*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

resmgr_context_t * resmgr_context_alloc
 (dispatch_t * dpp);
```

### **Arguments:**

*dpp*     A dispatch handle created by *dispatch\_create()*.

### **Library:**

`libc`

### **Description:**

The *resmgr\_context\_alloc()* function returns a context that's used for blocking and receiving messages.



---

This function is a special case of *dispatch\_context\_alloc()*. You should use it only when writing a simple resource manager.

---

### **Returns:**

A pointer to a `resmgr_context_t` structure, or NULL if an error occurs (*errno* is set).

### **Errors:**

EINVAL     No resource manager events were attached to *dpp*.  
ENOMEM    Insufficient memory to allocate the context *ctp*.

**Examples:**

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 dispatch_t *dpp;
 resmgr_context_t *ctp;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 :

 if ((ctp = resmgr_context_alloc (dpp)) == NULL) {
 fprintf(stderr, "Context wasn't allocated.\n");
 return EXIT_FAILURE;
 }
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*dispatch\_context\_alloc()*, *dispatch\_create()*, *resmgr\_attach()*,  
*resmgr\_context\_free()*, **resmgr\_context\_t**

“Components of a Resource Manager” in the Writing a Resource  
Manager chapter of the *Programmer’s Guide*

**Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

void resmgr_context_free(resmgr_context_t * ctp);
```

**Arguments:**

*ctp* A pointer to the `resmgr_context_t` structure that you want to free.

**Library:**

libc

**Description:**

The `resmgr_context_free()` function frees a context allocated by `resmgr_context_alloc()`.



---

This function is a special case of `dispatch_context_free()`. You should use it only when writing a simple resource manager.

---

**Examples:**

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 dispatch_t *dpp;
 resmgr_context_t *ctp;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }
 :
}
```

```
if ((ctp = resmgr_context_alloc (dpp)) == NULL) {
 fprintf(stderr, "Context wasn't allocated.\n");
 return EXIT_FAILURE;
}

:

resmgr_context_free (ctp);
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*resmgr\_context\_alloc()*, *resmgr\_context\_t*

“Components of a Resource Manager” in the Writing a Resource Manager chapter of the *Programmer's Guide*



**Synopsis:**

```
#include <sys/resmgr.h>

typedef struct _resmgr_context {
 int rcvid;
 struct _msg_info info;
 resmgr_iomsgs_t *msg;
 dispatch_t *dpp;
 int id;
 unsigned tid;
 unsigned msg_max_size;
 int status;
 int offset;
 int size;
 iov_t iov[1];
} resmgr_context_t;
```

**Description:**

The **resmgr\_context\_t** structure defines context information that's passed to resource-manager functions.

The members include:

|              |                                                                                                                         |
|--------------|-------------------------------------------------------------------------------------------------------------------------|
| <i>rcvid</i> | The receive ID to use for messages to and from the client.                                                              |
| <i>info</i>  | A pointer to a <b>_msg_info</b> structure that contains information about the message received by the resource manager. |
| <i>msg</i>   | A pointer to the message received by the resource manager, expressed as a union of all the possible message types.      |
| <i>dpp</i>   | The dispatch handle, created by <i>dispatch_create()</i> .                                                              |
| <i>id</i>    | The link Id, returned by <i>resmgr_attach()</i> .                                                                       |
| <i>tid</i>   | Not used; always zero.                                                                                                  |

|                     |                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------|
| <i>msg_max_size</i> | The minimum amount of space reserved for receiving a message.                                                |
| <i>status</i>       | A place to store the status of the current operation. Always use <i>_RESMGR_STATUS()</i> to set this member. |
| <i>offset</i>       | The offset, in bytes, into the client's message. You'll use this when working with combine messages.         |
| <i>size</i>         | The number of valid bytes in the message area.                                                               |
| <i>iov</i>          | An I/O vector where you can place the data that you're returning to the client.                              |

**Classification:**

QNX Neutrino

**See also:**

*dispatch\_create()*, *\_msg\_info*, *MsgInfo()*, *resmgr\_attach()*,  
*resmgr\_context\_alloc()*, *resmgr\_context\_free()*, *\_RESMGR\_STATUS()*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int resmgr_detach(dispatch_t * dpp,
 int id,
 unsigned flags);
```

### **Arguments:**

- dpp*      A dispatch handle created by *dispatch\_create()*.
- id*        The link ID that *resmgr\_attach()* returned.
- flags*     Flags that affect the operation. The possible *flags* (defined in *<sys/dispatch.h>* and *<sys/resmgr.h>*) are:
- *\_RESMGR\_DETACH\_ALL* — detach the name from the namespace and invalidate all open bindings.
  - *\_RESMGR\_DETACH\_PATHNAME* — detach only the name from the namespace, leaving existing bindings intact. This option is useful when you're unlinking a file or device, and you want to remove the name, but you want processes with open files to continue to use it until they close.

### **Library:**

`libc`

### **Description:**

The *resmgr\_detach()* function removes pathname *id* from the pathname space of context *dpp*.

## Blocking states

The *resmgr\_detach()* function blocks until the **RESMGR\_HANDLE\_T**, that's passed to the corresponding *resmgr\_attach()*, isn't being used in any connection function.

The effect that this has on servers is generally minimal. You should follow the following precautions to prevent potential deadlock situations:

- If you're using the **RESMGR\_HANDLE\_T** as an attribute, and that attribute is locked in any of the connection callouts (i.e. open, unlink, mount, etc.), then should unlock it before calling *resmgr\_detach()*. This allows any pending connection requests to complete before they're consequently invalidated.



---

If you call *resmgr\_detach()* from within a connection function, then the internal reference counting takes this into account and the server doesn't deadlock.

---

- If two or more *resmgr\_detach()* requests come in simultaneously, only one of the requests is served. The superfluous request will return with an error of -1 and *errno* set to ENOENT to indicate that the detachment process has already begun, and the entry is now invalid. If dynamically allocated, you should release **RESMGR\_HANDLE\_T** only after a successful return from *resmgr\_detach()*.
- If *resmgr\_detach()* is called and an existing client connection is established, then the I/O callout table is redirected for that client connection. The client will receive an error of EBADF when it uses the *fd* associated with that connection.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EINVAL      The *id* was never attached with *resmgr\_attach()*.
- ENOENT      A previous detachment request is in progress, or the *id* has already been detached.

**Examples:**

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 dispatch_t *dpp;
 int id;

 if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 id = resmgr_attach (...);

 :

 if (resmgr_detach(dpp, id, 0) == -1) {
 fprintf(stderr, "Failed to remove pathname \
 from the pathname space.\n");
 return EXIT_FAILURE;
 }
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

**Classification:**

QNX Neutrino

**Safety**

Cancellation point    Yes

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*dispatch\_create()*, *resmgr\_attach()*

“Writing a Resource Manager” chapter of the *Programmer’s Guide*.

## Synopsis:

```
#include <sys/resmgr.h>

int resmgr_devino(int id,
 dev_t *pdevno,
 ino64_t *pino);
```

## Arguments:

*id*            The link ID that *resmgr\_attach()* returned.

*pdevno*        A pointer to a **dev\_t** object where the function can store the device number.

*pino*           A pointer to a **ino64\_t** object where the function can store the inode number.

## Library:

**libc**

## Description:

The function *resmgr\_devino()* fills in the structures pointed to by *pdevno* and *pino* with the device number and inode number extracted from *id*.

This function is typically used to fill in:

- *iofunc\_mount\_t->dev*
- *iofunc\_attr\_t->inode*

You can use the *major()*, *minor()*, and *makedev()* macros to work with device IDs. They're defined in **<sys/types.h>** and are described in the documentation for *stat()*.

## Returns:

-1 on error (*errno* is set); any other value on success.

## Errors:

EINVAL The *id* argument is invalid.

## Examples:

```
#include <sys/resmgr.h>
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 iofunc_mount_t mount;
 iofunc_attr_t attr;

 ...
 attr.mount = &mount;
 ...
 id = resmgr_attach(...)
 ...
 resmgr_devino(id, &mount.dev, &attr.inode);
 ...
 return EXIT_SUCCESS;
}
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



**See also:**

*resmgr\_attach()*, *SETIOV()*, *stat()*

## ***resmgr\_handle\_tune()***

© 2004, QNX Software Systems Ltd.

*Tune aspects of client fd-to-OCB mapping*

### **Synopsis:**

```
int resmgr_handle_tune(int min_handles,
 int min_clients,
 int max_client_handles,
 int *min_handles_old,
 int *min_clients_old,
 int *max_client_handles_old);
```

### **Arguments:**

*min\_handles* To perform the described mapping, the resource manager framework makes use of **`_resmgr_handle_entry`** structures. This value describes the minimum number of these structures to keep around. If more than this number are in use, they may be returned to the heap via *free()* as they're released.

*min\_clients, max\_client\_handles*

To perform the described mapping, the resource manager framework makes use of hash buckets, one per client. The *min\_clients* describes the minimum number of these buckets to keep around. If more than this number of clients are in communication with your resource manager, these buckets may be released back to the heap via *free()* as particular clients close all their *fds* to your manager.

The *max\_client\_handles* describes the size of each of these hash buckets. The maximum number of lookups to find a particular *fd*-to-OCB mapping is the client's max *fd* divided by *max\_client\_handles* rounded to the nearest integer, i.e. in pseudocode:

**`ceil(max fd/max_client_handles)`**.

If this value changes, the new value takes effect for newly connected clients. Existing clients are unaffected.

If negative values are specified to any of the above three parameters, their current values are left unchanged.

*\*\_old*

If any of these are non-NULL, the corresponding value in use by the resource manager layer at the time of the call is returned.

## Library:

`libc`

## Description:

One of the functions of the resource manager framework is to perform the mapping of client file descriptors to structures local to the resource manager that describe these descriptors. These structures are often Open Control Blocks (OCBs). For details on OCBs, see *resmgr\_open\_bind()*. The *resmgr\_handle\_tune()* function can be used to tune certain aspects of this mapping and subsequent lookups of a client's OCBs.

## Returns:

0.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*resmgr\_open\_bind()*

Writing a Resource Manager chapter of the *Programmer's Guide*

**Synopsis:**

```
#include <resmgr.h>

int _resmgr_handle_grow (unsigned min);
```

**Arguments:**

*min*     The number of requests that you want to accommodate.

**Library:**

`libc`

**Description:**

The *\_resmgr\_handle\_grow()* function pre-grows or allocates the resource manager database table entries to support a given number of connections to improve runtime performance by reducing the number of dynamic memory allocations.

The function pre-allocates database space for *min* requests.

**Returns:**

The number of free entries in the table, or -1 if the resource manager table can't be locked.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

*\_resmgr\_handle\_grow()*

© 2004, QNX Software Systems Ltd.

## Synopsis:

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int resmgr_handler(resmgr_context_t * ctp);
```

## Arguments:

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

## Library:

`libc`

## Description:

The `resmgr_handler()` function handles the message received in *ctp*. This function handles different I/O messages through the resource manager framework.



---

The `resmgr_handler()` function is a special case of `dispatch_handler()`. You should use it only when writing a simple resource manager i.e. where there's no need to attach pulses or messages.

---

## Returns:

0 Success.  
-1 An error occurred.

## Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
```

```
dispatch_t *dpp;
resmgr_context_t *ctp;

if ((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
}

:

ctp = resmgr_context_alloc(dpp);

while (1) {
 if ((ctp = resmgr_block(ctp)) == NULL) {
 fprintf(stderr, "block error\n");
 return EXIT_FAILURE;
 }
 resmgr_handler(ctp);
}
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



## **Caveats:**

If you attach messages or pulses to *dpp* by calling *message\_attach()*, *pulse\_attach()*, or *select\_attach()*, those events aren't dispatched by *resmgr\_handler()*. Instead, you should call *dispatch\_handler()*.

## **See also:**

*dispatch\_handler()*, *resmgr\_attach()*, *resmgr\_block()*

“Components of a Resource Manager” in the Writing a Resource Manager chapter of the *Programmer's Guide*

## **\_resmgr\_io\_func()**

© 2004, QNX Software Systems Ltd.

*Retrieve an I/O function from an I/O function table*

### **Synopsis:**

```
#include <sys/resmgr.h>

_resmgr_func_t *_resmgr_io_func(
 const resmgr_io_funcs_t * funcs,
 unsigned type);
```

### **Arguments:**

*funcs* A pointer to the `resmgr_io_funcs_t` structure for the table of I/O functions.

*type* The type of I/O function that you want to get from the table. This argument should be one of the values defined in `<sys/iomsg.h>`, such as `_IO_READ` or `_IO_WRITE`.

### **Library:**

`libc`

### **Description:**

The `_resmgr_io_func()` function retrieves the I/O function associated with *type* from the function table defined by *funcs*.

### **Returns:**

A pointer to the function responsible for servicing *type*, or NULL if the function can't be found.

### **Classification:**

QNX Neutrino

#### **Safety**

---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*resmgr\_io\_funcs\_t*, *resmgr\_iofuncs()*

## resmgr\_io\_funcs\_t

© 2004, QNX Software Systems Ltd.

*Table of POSIX-level I/O functions*

### Synopsis:

```
#include <sys/resmgr.h>

typedef struct _resmgr_io_funcs {
 unsigned nfuncs;
 int (*read) (resmgr_context_t *ctp, io_read_t *msg,
 RESMGR_OCB_T *ocb);
 int (*write) (resmgr_context_t *ctp, io_write_t *msg,
 RESMGR_OCB_T *ocb);
 int (*close_ocb) (resmgr_context_t *ctp, void *reserved,
 RESMGR_OCB_T *ocb);
 int (*stat) (resmgr_context_t *ctp, io_stat_t *msg,
 RESMGR_OCB_T *ocb);
 int (*notify) (resmgr_context_t *ctp, io_notify_t *msg,
 RESMGR_OCB_T *ocb);
 int (*devctl) (resmgr_context_t *ctp, io_devctl_t *msg,
 RESMGR_OCB_T *ocb);
 int (*unlock) (resmgr_context_t *ctp, io_pulse_t *msg,
 RESMGR_OCB_T *ocb);
 int (*pathconf) (resmgr_context_t *ctp, io_pathconf_t *msg,
 RESMGR_OCB_T *ocb);
 int (*lseek) (resmgr_context_t *ctp, io_lseek_t *msg,
 RESMGR_OCB_T *ocb);
 int (*chmod) (resmgr_context_t *ctp, io_chmod_t *msg,
 RESMGR_OCB_T *ocb);
 int (*chown) (resmgr_context_t *ctp, io_chown_t *msg,
 RESMGR_OCB_T *ocb);
 int (*utime) (resmgr_context_t *ctp, io_utime_t *msg,
 RESMGR_OCB_T *ocb);
 int (*fdopen) (resmgr_context_t *ctp, io_fdopen_t *msg,
 RESMGR_OCB_T *ocb);
 int (*fdinfo) (resmgr_context_t *ctp, io_fdinfo_t *msg,
 RESMGR_OCB_T *ocb);
 int (*lock) (resmgr_context_t *ctp, io_lock_t *msg,
 RESMGR_OCB_T *ocb);
 int (*space) (resmgr_context_t *ctp, io_space_t *msg,
 RESMGR_OCB_T *ocb);
 int (*shutdown) (resmgr_context_t *ctp, io_shutdown_t *msg,
 RESMGR_OCB_T *ocb);
 int (*mmap) (resmgr_context_t *ctp, io_mmap_t *msg,
 RESMGR_OCB_T *ocb);
 int (*msg) (resmgr_context_t *ctp, io_msg_t *msg,
```

```

 RESMGR_OCB_T *ocb);
 int (*reserved) (resmgr_context_t *ctp, void *msg,
 RESMGR_OCB_T *ocb);
 int (*dup) (resmgr_context_t *ctp, io_dup_t *msg,
 RESMGR_OCB_T *ocb);
 int (*close_dup) (resmgr_context_t *ctp, io_close_t *msg,
 RESMGR_OCB_T *ocb);
 int (*lock_ocb) (resmgr_context_t *ctp, void *reserved,
 RESMGR_OCB_T *ocb);
 int (*unlock_ocb) (resmgr_context_t *ctp, void *reserved,
 RESMGR_OCB_T *ocb);
 int (*sync) (resmgr_context_t *ctp, io_sync_t *msg,
 RESMGR_OCB_T *ocb);
} resmgr_io_funcs_t;

```

## Description:

The `resmgr_connect_funcs_t` structure is a table of the POSIX-level I/O functions that are used by a resource manager. You can initialize this table by calling `iofunc_func_init()` and then overriding the defaults with your own functions.

This structure includes `nfuncs`, which indicates how many functions are in the table (in case the structure grows in the future), along with these functions:

| Member:      | Used to:                        | Default:                      |
|--------------|---------------------------------|-------------------------------|
| <i>read</i>  | Handle<br>_IO_READ<br>messages  | <i>iofunc_read_default()</i>  |
| <i>write</i> | Handle<br>_IO_WRITE<br>messages | <i>iofunc_write_default()</i> |

*continued...*

| <b>Member:</b>   | <b>Used to:</b>                                 | <b>Default:</b>                   |
|------------------|-------------------------------------------------|-----------------------------------|
| <i>close_ocb</i> | Return the memory allocated for an OCB          | <i>iofunc_close_ocb_default()</i> |
| <i>stat</i>      | Handle _IO_STAT messages                        | <i>iofunc_stat_default()</i>      |
| <i>notify</i>    | Handle _IO_NOTIFY messages                      | None                              |
| <i>devctl</i>    | Handle _IO_DEVCTL messages                      | <i>iofunc_devctl_default()</i>    |
| <i>unblock</i>   | Unblock the resource if an operation is aborted | <i>iofunc_unblock_default()</i>   |
| <i>pathconf</i>  | Handle _IO_PATHCONF messages                    | <i>iofunc_pathconf_default()</i>  |
| <i>lseek</i>     | Handle _IO_LSEEK messages                       | <i>iofunc_lseek_default()</i>     |
| <i>chmod</i>     | Handle _IO_CHMOD messages                       | <i>iofunc_chmod_default()</i>     |
| <i>chown</i>     | Handle _IO_CHOWN messages                       | <i>iofunc_chown_default()</i>     |

*continued...*

| <b>Member:</b>  | <b>Used to:</b>                                                             | <b>Default:</b>                  |
|-----------------|-----------------------------------------------------------------------------|----------------------------------|
| <i>utime</i>    | Handle<br>_IO_UTIME<br>messages                                             | <i>iofunc_utime_default()</i>    |
| <i>fdopen</i>   | Handle<br>_IO_OPENFD<br>messages                                            | <i>iofunc_openfd_default()</i>   |
| <i>fdinfo</i>   | Handle<br>_IO_FDINFO<br>messages                                            | <i>iofunc_fdinfo_default()</i>   |
| <i>lock</i>     | Handle<br>_IO_LOCK<br>messages                                              | <i>iofunc_lock_default()</i>     |
| <i>space</i>    | Allocate or free<br>memory for the<br>resource.                             | None                             |
| <i>shutdown</i> | Reserved                                                                    | None                             |
| <i>mmap</i>     | Handle<br>_IO_MMAP<br>messages                                              | <i>iofunc_mmap_default()</i>     |
| <i>msg</i>      | Handle a more<br>general<br>messaging<br>scheme to<br>control the<br>device | None                             |
| <i>reserved</i> | Not applicable                                                              | None                             |
| <i>dup</i>      | Handle<br>_IO_DUP<br>messages                                               | None — handled by the base layer |

*continued...*

| Member:           | Used to:                                               | Default:                           |
|-------------------|--------------------------------------------------------|------------------------------------|
| <i>close_dup</i>  | Handle<br>_IO_CLOSE<br>messages                        | <i>iofunc_close_dup_default()</i>  |
| <i>lock_ocb</i>   | Lock the<br>attributes for a<br>group of<br>messages   | <i>iofunc_lock_ocb_default()</i>   |
| <i>unlock_ocb</i> | Unlock the<br>attributes for a<br>group of<br>messages | <i>iofunc_unlock_ocb_default()</i> |
| <i>sync</i>       | Handle<br>_IO_SYNC<br>messages                         | <i>iofunc_sync_default()</i>       |



If you use *iofunc\_lock\_default()*, you must also use *iofunc\_close\_dup\_default()* and *iofunc\_unlock\_default()*.

## Classification:

QNX Neutrino

## See also:

*iofunc\_chmod\_default()*, *iofunc\_chown\_default()*,  
*iofunc\_close\_dup\_default()*, *iofunc\_close\_ocb\_default()*,  
*iofunc\_devctl\_default()*, *iofunc\_fdinfo\_default()*, *iofunc\_func\_init()*,  
*iofunc\_lock\_default()*, *iofunc\_lock\_ocb\_default()*,  
*iofunc\_lseek\_default()*, *iofunc\_mmap\_default()*,  
*iofunc\_openfd\_default()*, *iofunc\_pathconf\_default()*,  
*iofunc\_read\_default()*, *iofunc\_stat\_default()*, *iofunc\_sync\_default()*,  
*iofunc\_unblock\_default()*, *iofunc\_unlock\_ocb\_default()*,  
*iofunc\_utime\_default()*, *iofunc\_write\_default()*,  
**resmgr\_connect\_funcs\_t**



Writing a Resource Manager chapter of the QNX Neutrino  
*Programmer's Guide*

## ***resmgr\_iofuncs()***

© 2004, QNX Software Systems Ltd.

*Extract the I/O function pointers associated with client connections*

### **Synopsis:**

```
#include <resmgr.h>

const resmgr_io_funcs_t * resmgr_iofuncs(
 resmgr_context_t * ctp;
 struct _msg_info * info);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*info* A pointer to the `_msg_info` structure that describes the binding to the client. You should fill this structure by calling `MsgInfo()`.

### **Library:**

`libc`

### **Description:**

The `resmgr_iofuncs()` function retrieves the I/O function callout table associated with the client connections described by binding specified by *info*.

Note that context information pointed to by *ctp* actually contains *info*.

### **Returns:**

A pointer to the `resmgr_io_funcs_t` I/O function callout table, or NULL if the binding can't be found or an error occurs.

### **Errors:**

ESRCH The connection can't be located in the resource manager's table.

ENOMEM There is no memory available for the operation.

EINVAL      Invalid arguments were used.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

`_msg_info`, `MsgInfo()`, `resmgr_io_funcs_t`, `_resmgr_ocb()`,  
`resmgr_open_bind()`, `resmgr_unbind()`

## ***resmgr\_msgread()***

© 2004, QNX Software Systems Ltd.

*Read a message from a client*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_msgread(resmgr_context_t * ctp,
 void * msg,
 int size,
 int offset);
```

### **Arguments:**

|               |                                                                                                                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>    | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions. This function extracts the <code>rcvid</code> from this structure. |
| <i>msg</i>    | A pointer to a buffer where the function can store the data.                                                                                                                                                    |
| <i>bytes</i>  | The number of bytes that you want to read. These functions don't let you read past the end of the thread's message; they return the number of bytes actually read.                                              |
| <i>offset</i> | An offset into the thread's send message that indicates where you want to start reading the data.                                                                                                               |

### **Library:**

`libc`

### **Description:**

The `resmgr_msgread()` function is a convenience function that you should in a resource manager instead of `MsgRead()`.

You'll use `resmgr_msgread()` when you handle *combine messages*, where the offset of the rest of the message that's to be read is additionally offset by previous combine message elements. For more information, see "Combine messages" in the Writing a Resource Manager chapter of the *Programmer's Guide*.

**Returns:**

The same values as *MsgRead()*: the number of bytes read, or -1 if an error occurs (*errno* is set).

**Errors:**

|            |                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in a server's address space when it tried to access the caller's message buffers. |
| ESRCH      | The thread indicated by <i>ctp</i> -> <i>rcvid</i> doesn't exist or its connection is detached.    |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                             |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgRead()*, *resmgr\_context\_t*, *resmgr\_msgreadv()*,  
*resmgr\_msgwrite()*, *resmgr\_msgwritev()*

“Combine messages” in the Writing a Resource Manager chapter of the *Programmer's Guide*

## ***resmgr\_msgreadv()***

© 2004, QNX Software Systems Ltd.

*Read a message from a client*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_msgreadv(resmgr_context_t * ctp,
 iov_t * rmsg,
 int rparts,
 int offset);
```

### **Arguments:**

|               |                                                                                                                                                                                                     |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>    | A pointer to a <b>resmgr_context_t</b> structure that the resource-manager library uses to pass context information between functions. This function extracts the <i>rcvid</i> from this structure. |
| <i>riov</i>   | An array of buffers where the functions can store the data.                                                                                                                                         |
| <i>rparts</i> | The number of elements in the <i>riov</i> array.                                                                                                                                                    |
| <i>offset</i> | An offset into the thread's send message that indicates where you want to start reading the data.                                                                                                   |

### **Library:**

**libc**

### **Description:**

This *resmgr\_msgreadv()* function is a convenience function that you should use in a resource manager instead of *MsgReadv()*.

You'll use *resmgr\_msgreadv()* when handling combine messages, where the offset of the rest of the message that is to be read is additionally offset by previous combine message elements. For more information, see "Combine messages" in the Writing a Resource Manager chapter of the *Programmer's Guide*.

**Returns:**

The same values as *MsgReadv()*: the number of bytes read, or -1 if an error occurs (*errno* is set).

**Errors:**

|            |                                                                                                      |
|------------|------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in a server's address space when it tried to access the caller's message buffers.   |
| ESRCH      | The thread indicated by <i>ctp</i> -> <i>rcvid</i> doesn't exist or has had its connection detached. |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                               |

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgReadv()*, *resmgr\_context\_t*, *resmgr\_msgread()*,  
*resmgr\_msgwrite()*, *resmgr\_msgwritev()*

“Combine messages” in the Writing a Resource Manager chapter of the *Programmer's Guide*

## ***resmgr\_msgwrite()***

© 2004, QNX Software Systems Ltd.

*Write a message to a client*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_msgwrite(resmgr_context_t *ctp,
 const void *msg,
 int size,
 int offset);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions. This function extracts the `rcvid` from this structure.

*msg* A pointer to a buffer that contains the data you want to write.

*size* The number of bytes that you want to write. These functions don't let you write past the end of the sender's buffer; they return the number of bytes actually written.

*offset* An offset into the sender's buffer that indicates where you want to start writing the data.

### **Library:**

`libc`

### **Description:**

The function `resmgr_msgwrite()` is a cover for `MsgWrite()` and performs the exact same functionality.

### **Returns:**

The same values as `MsgWrite()`; the number of bytes written, or -1 if an error occurs (`errno` is set).



**Errors:**

|            |                                                                                                                  |
|------------|------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffer. |
| ESRCH      | The thread indicated by <i>ctp</i> -> <i>rvid</i> does not exist or has had its connection detached.             |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                           |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgWrite()*, **resmgr\_context\_t**, *resmgr\_msgread()*,  
*resmgr\_msgreadv()*, *resmgr\_msgwritev()*

“Combine messages” in the Writing a Resource Manager chapter of the *Programmer's Guide*

## ***resmgr\_msgwritev()***

© 2004, QNX Software Systems Ltd.

*Write a message to a client*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_msgwritev(resmgr_context_t *ctp,
 const iov_t *msg,
 int sparts,
 int offset);
```

### **Arguments:**

*ctp*      A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions. This function extracts the rclid from this structure.

*iov*      An array of buffers that contains the data you want to write.

*parts*    The number of elements in the array. These functions don't let you write past the end of the sender's buffer; they return the number of bytes actually written.

*offset*   An offset into the sender's buffer that indicates where you want to start writing the data.

### **Library:**

libc

### **Description:**

The *resmgr\_msgwritev()* function is a cover function for *MsgWritev()*, and performs the exact same functionality. It's provided for consistency with *resmgr\_msgwrite()*.

### **Returns:**

The number of bytes written, or -1 if an error occurred (*errno* is set).

**Errors:**

|            |                                                                                                                  |
|------------|------------------------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in the sender's address space when a server tried to access the sender's return message buffer. |
| ESRCH      | The thread indicated by <i>ctp</i> -> <i>rcvid</i> does not exist or has had its connection detached.            |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                                           |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*MsgWritev()*, **resmgr\_context\_t**, *resmgr\_msgread()*,  
*resmgr\_msgreadv()*, *resmgr\_msgwrite()*

“Combine messages” in the Writing a Resource Manager chapter of the *Programmer's Guide*

## **\_RESMGR\_NPARTS()**

© 2004, QNX Software Systems Ltd.

*Get a given number of parts from the `ctp->iov` structure*

### **Synopsis:**

```
#include <sys/resmgr.h>

_RESMGR_NPARTS(int num)
```

### **Arguments:**

*num*      The number of parts that you want to get.

### **Library:**

`libc`

### **Description:**

The macro `_RESMGR_NPARTS()` indicates to the caller to get *num* parts from the `ctp->iov` structure (see `resmgr_context_t`). The macro is similar to:

*MsgReply( ctp->rvid, ctp->status, ctp->iov, num ).*

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The *resmgr\_attach()* function should set *attr->nparts\_max* to be the maximum value for *num*.

**See also:**

*MsgReply()*, *resmgr\_attach()*, *resmgr\_context\_t*,  
*RESMGR\_PTR()*, *RESMGR\_STATUS()*

## **\_resmgr\_ocb()**

© 2004, QNX Software Systems Ltd.

*Retrieve an Open Control Block*

### **Synopsis:**

```
#include <sys/resmgr.h>

void * _resmgr_ocb(resmgr_context_t * ctp,
 struct _msg_info * info);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

Note that *ctp* actually contains *info*.

*info* A pointer to a `_msg_info` structure that indicates which client mapping you want to retrieve information about. To fill in the structure, call `MsgInfo()`.

### **Library:**

`libc`

### **Description:**

The `_resmgr_ocb()` function queries the internal resource manager database, which maps client connections to the server Open Control Block (OCB), to retrieve the OCB pointer that was previously bound using `resmgr_open_bind()`.

### **Returns:**

A pointer to the OCB for the matching binding, or NULL if the binding can't be found or an error occurred.

The OCB can be a structure that you define. By default, it's of type `iofunc_ocb_t`.

## **Errors:**

|        |                                                                |
|--------|----------------------------------------------------------------|
| ESRCH  | The connection can't be located in the resource manager table. |
| ENOMEM | There isn't enough memory available for the operation.         |
| EINVAL | Invalid arguments were used.                                   |

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## **See also:**

*iofunc\_ocb\_t*, *\_msg\_info*, *MsgInfo()*, *resmgr\_iofuncs()*, *resmgr\_open\_bind()*, *resmgr\_unbind()*

## ***resmgr\_open\_bind()***

© 2004, QNX Software Systems Ltd.

*Associate an OCB with an open request*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_open_bind(
 resmgr_context_t* ctp,
 void* ocb,
 const resmgr_io_funcs_t* iofuncs);
```

### **Arguments:**

*ctp* A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*ocb* A pointer to the Open Control Block that you want to bind to the open request. An OCB is usually a structure of type `iofunc_ocb_t`, but you can define your own.

*iofuncs* A pointer to the `resmgr_io_funcs_t` structure that defines the I/O functions for the resource manager.

### **Library:**

`libc`

### **Description:**

The `resmgr_open_bind()` function is the lowest-level call in the resource manager library used for handling open messages. It associates the Open Control Block (OCB) with a process identified by the `id` and `info` members of `ctp`.





---

You must use this function as part of the handling of an `_IO_OPEN` message. In practice, you don't call this function directly; you typically use either `iofunc_open_default()` or `iofunc_ocb_attach()`. (The `iofunc_open_default()` function calls `iofunc_ocb_attach()`, which in turn calls `resmgr_open_bind()`).

---

An internal data structure is allocated that maintains the number of links to the OCB. On a file descriptor `dup()`, the link count is incremented and on a `close()` it's decremented. When the count reaches zero, the `close_ocb()` callout specified in `io_funcs` is called.

In the most general case, the OCB is an arbitrary structure that you define that can hold information describing an open file, or just a simple `int` to hold the open mode for checking in the `read()` and `write()` callouts.

In the typical case, however, the OCB is a structure that contains at least the members as defined by the typedef `iofunc_ocb_t`. This typedef defines a common OCB structure that can then be used by the POSIX layer helper functions (all functions beginning with the name `iofunc_*`). The advantage of this approach is that your resource manager gets POSIX behavior for free, without any additional work on your part.

The `attr` argument to the `open()` callout is also typically saved in the OCB. The well defined `iofunc_ocb_t` has a member called `attr` to which you must assign the value of the `attr` argument. This lets the POSIX helper functions access information about the current open session (as stored in the OCB) as well as information about the device itself (as stored in the attributes structure, `ocb -> attr`).

For a detailed discussion, including several examples, see the Writing a Resource Manager chapter of the *Programmer's Guide*.

## Returns:

- 0      Success.
- 1     An error occurred (`errno` is set).

**Errors:**

- EINVAL      The *id* and/or *info* members of *ctp* aren't valid.
- ENOMEM     Insufficient memory to allocate an internal data structure.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        No
- Thread                 Yes

**See also:**

*iofunc\_ocb\_attach()*, *iofunc\_ocb\_t*, *iofunc\_open\_default()*,  
*resmgr\_context\_t*, *resmgr\_io\_funcs\_t*, *resmgr\_unbind()*

Writing a Resource Manager chapter of the *Programmer's Guide*

## ***resmgr\_pathname()***

*Return the pathname associated with an ID*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_pathname(int id,
 unsigned flags,
 char* path,
 int maxbuf);
```

### **Arguments:**

*id*            The link ID that *resmgr\_attach()* returned.

*flags*        Flags that affect the operation:

- `_RESMGR_PATHNAME_LOCALPATH` — get a shortened pathname that's usable only on the local machine. By default, this function gets a globally unique pathname.

*path*        A pointer to a buffer where the function can store the path name.

*maxbuf*      The size of the buffer.

### **Library:**

`libc`

### **Description:**

The *resmgr\_pathname()* function returns the pathname associated with an *id* that's returned from *resmgr\_attach()*, it's also the *ctp->id* value of all the resmgr callout functions.

If the *id* was obtained from calling *resmgr\_attach()* with `_RESMGR_FLAG_DIR` specified, then the path name includes a trailing slash.

By default, this function calls:

```
netmgr_ndtostr(ND2S_DIR_SHOW, nd, buf, sizeofbuf)
```

If you specify `_RESMGR_PATHNAME_LOCALPATH`, it calls

```
netmgr_ndtostr(ND2S_DIR_SHOW|ND2S_LOCAL_STR, nd, buf, sizeofbuf)
```

to return a shortened path that's usable on your local node only. This is useful for display.

## Returns:

The length of the path, including the terminating NULL character, or -1 if an error occurs (*errno* is set).

## Errors:

|            |                                                                                                    |
|------------|----------------------------------------------------------------------------------------------------|
| EFAULT     | A fault occurred in a server's address space when it tried to access the caller's message buffers. |
| ESRCH      | The thread indicated by <i>ctp</i> -> <i>rcvid</i> doesn't exist or its connection is detached.    |
| ESRVRFAULT | A fault occurred when the kernel tried to access the buffers provided.                             |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*netmgr\_ndtostr(), resmgr\_attach()*

## **\_RESMGR\_PTR()**

© 2004, QNX Software Systems Ltd.

*Get one part from the `ctp->iov` structure and fill in its fields*

### **Synopsis:**

```
#include <sys/resmgr.h>

_RESMGR_PTR(resmgr_context_t ctp,
 void msg,
 size_t nbytes)
```

### **Arguments:**

|               |                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>    | A pointer to a <code>resmgr_context_t</code> structure that the resource-manager library uses to pass context information between functions. |
| <i>msg</i>    | The value you want to use for the structure's <code>iov_base</code> member.                                                                  |
| <i>nbytes</i> | The value you want to use for the structure's <code>iov_len</code> member.                                                                   |

### **Library:**

`libc`

### **Description:**

The `_RESMGR_PTR()` macro gets one part from the `ctp->iov` structure (see `resmgr_context_t`) and fills in its fields. The macro is equivalent to:

`SETIOV(ctp->iov, msg, nbytes)`

returning `_RESMGR_NPARTS` (1).

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*resmgr\_context\_t*, *\_RESMGR\_NPARTS()*, *\_RESMGR\_STATUS()*,  
*SETIOV()*

## **RESMGR\_STATUS()**

© 2004, QNX Software Systems Ltd.

*Set the status member of a resource-manager context*

### **Synopsis:**

```
#include <sys/resmgr.h>

RESMGR_STATUS(resmgr_context_t *ctp,
 int status)
```

### **Arguments:**

*ctp*      A pointer to a `resmgr_context_t` structure that the resource-manager library uses to pass context information between functions.

*status*    The status that you want to set.

### **Library:**

`libc`

### **Description:**

The `RESMGR_STATUS()` macro sets the *status* member in the `resmgr_context_t` structure.

The resource manager library uses this status when it returns the value from `RESMGR_NPARTS()` for an I/O or connect function, such as:

`MsgReply ( ctp->rcvid, ctp->status, ctp->iov, num )`.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

`resmgr_connect_funcs_t`, `resmgr_io_funcs_t`,  
`_RESMGR_NPARTS()`, `_RESMGR_PTR()`

## ***resmgr\_unbind()***

© 2004, QNX Software Systems Ltd.

*Disassociate an OCB from an open request*

### **Synopsis:**

```
#include <sys/resmgr.h>

int resmgr_unbind(resmgr_context_t * ctp);
```

### **Arguments:**

*ctp* A pointer to a **resmgr\_context\_t** structure that the resource-manager library uses to pass context information between functions.

### **Library:**

**libc**

### **Description:**

The *resmgr\_unbind()* function removes a binding in the internal resources manager database (which maps client connections to server OCB pointers). The binding must previously have been bound by *resmgr\_open\_bind()*

The binding is reference counted; if multiple connections are established with the same binding, the binding is freed only when the last connection is removed.

You should use *MsgInfo()* to fill the **info** structure in *ctp* with information about which client mapping to retrieve.

### **Returns:**

-1 Failure.  
0 Success.

## Errors:

EINVAL      The binding can't be located in the resource manager table.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*MsgInfo()*, `resmgr_context_t`, *resmgr\_open\_bind()*

Writing a Resource Manager chapter of the *Programmer's Guide*

## ***rewind()***

© 2004, QNX Software Systems Ltd.

*Rewind a file stream to the beginning of the file*

### **Synopsis:**

```
#include <stdio.h>

void rewind(FILE *fp);
```

### **Arguments:**

*fp* The file stream that you want to rewind.

### **Library:**

libc

### **Description:**

The *rewind()* function rewinds the file stream specified by *fp* to the beginning of the file. It's equivalent to calling *fseek()* like this:

```
fseek(fp, 0L, SEEK_SET);
```

except that the error indicator for the stream is cleared.

### **Examples:**

This example shows how you might implement a two-pass assembler:

```
#include <stdio.h>
#include <stdlib.h>

void assemble_pass(FILE *fp, int passno)
{
 printf("Pass %d\n", passno);

 /* Do more work on the fp */
 switch(passno) {
 case 1:
 /* do the first-pass work */
 break;

 case 2:
 /* do the second-pass work */
 break;
 }
```

```
 default:
 break;
 }
 }

int main(void)
{
 FILE *fp;

 fp = fopen("program.s", "r");
 if(fp != NULL) {
 assemble_pass(fp, 1);
 rewind(fp);

 assemble_pass(fp, 2);
 fclose(fp);

 return EXIT_SUCCESS;
 }

 puts("Error opening program.s");

 return EXIT_FAILURE;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*clearerr(), fopen(), fseek()*

### **Synopsis:**

```
#include <sys/types.h>
#include <dirent.h>

void rewinddir(DIR * dirp);
```

### **Arguments:**

*dirp*     A pointer to the directory stream of the directory to rewind.

### **Library:**

`libc`

### **Description:**

The *rewinddir()* function rewinds the directory stream specified by *dirp* to the start of the directory. The directory stream will now refer to the current state of the directory, as if the calling thread had called *opendir()* again.



---

The result of using a directory stream after one of the *exec\*()* or *spawn\*()* family of functions is undefined. After a call to *fork()*, either the parent *or* the child (but not both) can continue processing the directory stream, using the *readdir()* and *rewinddir()* functions. If both the parent and child processes use these functions, the result is undefined. Either (or both) processes may use *closedir()*.

---

### **Examples:**

List all the files in a directory, create a new file, and then list the directory contents again:

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <dirent.h>
#include <stdlib.h>
```

```
int main(void)
{
 DIR *dirp;
 struct dirent *direntp;
 int filedes;

 dirp = opendir("/home/fred");
 if(dirp != NULL) {
 printf("Old directory listing\n");
 for(;;) {
 direntp = readdir(dirp);
 if(direntp == NULL) break;
 printf("%s\n", direntp->d_name);
 }

 filedes = creat("/home/fred/file.new",
 S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP);
 close(filedes);

 rewinddir(dirp);
 printf("New directory listing\n");
 for(;;) {
 direntp = readdir(dirp);
 if(direntp == NULL) break;
 printf("%s\n", direntp->d_name);
 }
 closedir(dirp);
 }

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |

*continued...*



**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*closedir(), opendir(), readdir(), readdir\_r(), seekdir()*

## ***Rgetsockname()***

© 2004, QNX Software Systems Ltd.

*Get the name of a socket (via a SOCKS server)*

### **Synopsis:**

```
#include <sys/socket.h>

int Rgetsockname(int s,
 struct sockaddr * name,
 int * namelen);
```

### **Arguments:**

|                |                                                                                                                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>       | The file descriptor of the socket whose name you want to get.                                                                                                                                          |
| <i>name</i>    | A pointer to a <b>sockaddr</b> object where the function can store the socket's name.                                                                                                                  |
| <i>namelen</i> | A pointer to a <b>socklen_t</b> object that initially indicates the amount of space pointed to by <i>name</i> . The function updates <i>namelen</i> to contain the actual size of the name (in bytes). |

### **Library:**

**libsocks**

### **Description:**

The *Rgetsockname()* function is a cover function for *getsockname()* — the difference is that *Rgetsockname()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

### **Returns:**

|    |                                           |
|----|-------------------------------------------|
| 0  | Success.                                  |
| -1 | An error occurred ( <i>errno</i> is set). |

**Classification:**

SOCKS

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*getsockname()* *Raccept()*, *Rbind()*, *Rconnect()*, *Rlisten()*, *Rrcmd()*,  
*Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall

## ***rindex()***

© 2004, QNX Software Systems Ltd.

*Find the last occurrence of a character in a string*

### **Synopsis:**

```
#include <strings.h>

char *rindex(const char *s,
 int c);
```

### **Arguments:**

- s*     The string you want to search. This string must end with a null (`\0`) character. The null character is considered to be part of the string.
- c*     The character you're looking for.

### **Library:**

`libc`

### **Description:**

The *rindex()* function returns a pointer to the last occurrence of the character *c* in the string *s*.

### **Returns:**

A pointer to the character, or NULL if the character doesn't occur in the string.

### **Classification:**

Legacy Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*index(), strchr(), strrchr()*

## ***rint()*, *rintf()***

© 2004, QNX Software Systems Ltd.

*Round to the nearest integral value*

---

### **Synopsis:**

```
#include <math.h>

double rint (double x);

float rintf (float x);
```

### **Arguments:**

*x*     The number that you want to round.

### **Library:**

`libm`

### **Description:**

The *rint()* and *rintf()* functions return the integral value nearest *x* in the direction of the current rounding mode.

If the current rounding mode rounds toward negative infinity, then *rint()* is identical to *floor()*. If the current rounding mode rounds toward positive infinity, then *rint()* is identical to *ceil()*.

### **Returns:**

An integer (represented as a double precision number) nearest *x* in the direction of the current rounding mode (*IEEE754*).

| <b>If <i>x</i> is:</b> | <b><i>rint()</i> returns:</b> |
|------------------------|-------------------------------|
|------------------------|-------------------------------|

|                |          |
|----------------|----------|
| $\pm$ Infinity | <i>x</i> |
|----------------|----------|

|     |     |
|-----|-----|
| NAN | NAN |
|-----|-----|



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

### Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b;

 a = 0.7 ;
 b = rint(a);
 printf("Round Native mode %f -> %f \n", a, b);

 return(0);
}
```

### Classification:

*rint()* is standard Unix; *rintf()* is ANSI (draft)

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ceil()*, *floor()*



### **Synopsis:**

```
#include <sys/socket.h>

int Rlisten(int s,
 int backlog);
```

### **Arguments:**

*s*            The descriptor for the socket that you want to listen on. You can create a socket by calling *socket()*.

*backlog*     The maximum length that the queue of pending connections may grow to.

### **Library:**

`libsocks`

### **Description:**

The *Rlisten()* function is a cover function for *listen()* — the difference is that *Rlisten()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

### **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

### **Classification:**

SOCKS

## **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## **See also:**

*listen()* *Raccept()*, *Rbind()*, *Rconnect()*, *Rgetsockname()*, *Rrcmd()*,  
*Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall

## Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

int rmdir(const char* path);
```

## Arguments:

*path* The path of the directory that you want to delete. This path can be relative to the current working directory, or an absolute path.

## Library:

`libc`

## Description:

The *rmdir()* function removes (deletes) the specified directory. The directory must not contain any files or directories.



If the directory is the current working directory of any process, *rmdir()* returns `-1` and sets *errno* to `EINVAL`. If the directory is the root directory, the effect of this function depends on the filesystem.

---

The space occupied by the directory is freed, making it inaccessible, if its link count becomes zero and no process has the directory open (*opendir()*). If a process has the directory open when the last link is removed, the `.` and `..` entries are removed and no new entries can be created in the directory. In this case, the directory will be removed when all references to it have been closed (*closedir()*).

When successful, *rmdir()* marks *st\_ctime* and *st\_mtime* for update in the parent directory.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EACCES Search permission is denied for a component of *path*, or write permission is denied on the parent directory of the directory to be removed.
- EBUSY The directory named by *path* can't be removed because it's being used by another process, and the implementation considers this to be an error.
- EEXIST The *path* argument names a directory that isn't empty.
- ELOOP Too many levels of symbolic links.
- ENAMETOOLONG  
The argument *path* exceeds PATH\_MAX in length, or a pathname component is longer than NAME\_MAX.
- ENOENT The specified *path* doesn't exist, or *path* is an empty string.
- ENOSYS The *rmdir()* function isn't implemented for the filesystem specified in *path*.
- ENOTDIR A component of *path* isn't a directory.
- ENOTEMPTY The *path* argument names a directory that isn't empty.
- EROFS The directory entry to be removed resides on a read-only filesystem.

**Examples:**

To remove the directory called `/home/terry`:

```
#include <stdlib.h>
#include <sys/types.h>
#include <sys/stat.h>

int main(void)
{
 (void)rmmdir("/home/terry");

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*chdir()*, *chmod()*, *errno*, *getcwd()*, *mkdir()*, *stat()*

# ROUTE

© 2004, QNX Software Systems Ltd.

*System packet forwarding database*

## Synopsis:

```
#include <sys/socket.h>
#include <net/if.h>
#include <net/route.h>

int socket(PF_ROUTE,
 SOCK_RAW,
 int family);
```

## Description:

QNX TCP/IP provides some packet routing facilities.



---

The following information applies only to the full TCP/IP stack. For information on how the tiny TCP/IP stack can change the routing table, see `npm-ttcpip.so`.

---

The socket manager maintains a routing information database that's used in selecting the appropriate network interface when transmitting packets.

A user process (or possibly multiple cooperating processes) maintains this database by sending messages over a special kind of socket. This supplants fixed-size *ioctl()*s used in earlier releases. Routing table changes may be carried out only by the superuser.

This interface may spontaneously emit routing messages in response to external events, such as receipt of a redirect or of a failure to locate a suitable route for a request. The message types are described in greater detail below.

## Routing database entries

Routing database entries come in two flavors: for a specific host or for all hosts on a generic subnetwork (as specified by a bit mask and value under the mask). The effect of wildcard or default routing may be achieved by using a mask of all zeros. There may be hierarchical routes.

When the system is booted and addresses are assigned to the network interfaces, each protocol family installs a routing table entry for each interface when it's ready for traffic. Normally the protocol specifies the route through each interface as a "direct" connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface is requested to address the packet to the gateway listed in the routing entry (i.e. the packet is forwarded).

### Routing packets

When routing a packet, the kernel attempts to find the most specific route matching the destination. (If there are two different mask and value-under-the-mask pairs that match, the more specific is the one with more bits in the mask. A route to a host is regarded as being supplied with a mask of as many ones as there are bits in the destination). If no entry is found, the destination is declared to be unreachable, and a routing-miss message is generated if there are any listeners on the routing control socket described below.

A wildcard routing entry is specified with a zero destination address value and a mask of all zeroes. Wildcard routes are used when the system fails to find other routes matching the destination. The combination of wildcard routes and routing redirects can provide an economical mechanism for routing traffic.

### Routing control messages

To open the channel for passing routing control messages, use the socket call shown in the synopsis above.

The family parameter may be `AF_UNSPEC`, which provides routing information for all address families, or it can be restricted to a specific address family. There can be more than one routing socket open per system.

Messages are formed by a header followed by a small number of `sockaddrs` (with variable length) interpreted by position and delimited by the new length entry in the `sockaddr`. An example of a

message with four addresses might be a redirect: Destination, Netmask, Gateway, and Author of the redirect. The interpretation of which addresses are present is given by a bit mask within the header; the sequence is least-significant to most-significant bit within the vector.

Any messages sent to the socket manager are returned, and copies are sent to all interested listeners. The interface provides the process ID for the sender. To distinguish between outstanding messages, the sender may use an additional sequence field. However, message replies may be lost when socket manager buffers are exhausted.

The interface may reject certain messages, as indicated in the *rtm\_errno* field.

| <b>This error occurs:</b> | <b>If:</b> |
|---------------------------|------------|
|---------------------------|------------|

|         |                                                               |
|---------|---------------------------------------------------------------|
| EEXIST  | Requested to duplicate an existing entry.                     |
| ESRCH   | Requested to delete a nonexistent entry.                      |
| ENOBUFS | Insufficient resources were available to install a new route. |

In the current implementation, all routing processes run locally, and the values for *rtm\_errno* are available through the normal *errno* mechanism, even if the routing reply message is lost.

A process may avoid the expense of reading replies to its own messages by calling *setsockopt()*, to turn off the SO\_USELOOPBACK option at the SOL\_SOCKET level. A process may ignore all messages from the routing socket by doing a *shutdown()* system call for further input.

If a route is in use when it's deleted, the routing entry is marked down and removed from the routing table, but the resources associated with it won't be reclaimed until all references to it are released. User processes can obtain information about the routing entry to a specific destination by using a RTM\_GET message or by calling *sysctl()*.



The messages are:

```
#define RTM_ADD 0x1 /* Add Route */
#define RTM_DELETE 0x2 /* Delete Route */
#define RTM_CHANGE 0x3 /* Change Metrics, Flags, or Gateway */
#define RTM_GET 0x4 /* Report Information */
#define RTM_LOSING 0x5 /* Kernel Suspects Partitioning */
#define RTM_REDIRECT 0x6 /* Told to use different route */
#define RTM_MISS 0x7 /* Lookup failed on this address */
#define RTM_RESOLVE 0xb /* request to resolve dst to LL addr */
#define RTM_NEWADDR 0xc /* address being added to iface */
#define RTM_DELADDR 0xd /* address being removed from iface */
#define RTM_IFINFO 0xe /* iface going up/down etc. */
```

A message header consists of one of the following:

```
struct rt_msghdr {
 u_short rtm_msglen; /* skip over non-understood msgs */
 u_char rtm_version; /* future binary compatibility */
 u_char rtm_type; /* message type */
 u_short rtm_index; /* index for associated ifp */
 int rtm_flags; /* flags, incl kern & message, e.g. DONE */
 int rtm_addr; /* bitmask identifying sockaddrs in msg */
 pid_t rtm_pid; /* identify sender */
 int rtm_seq; /* for sender to identify action */
 int rtm_errno; /* why failed */
 int rtm_use; /* from rtentry */
 u_long rtm_inits; /* which metrics we're initializing */
 struct rt_metrics rtm_rmx; /* metrics themselves */
};

struct if_msghdr {
 u_short ifm_msglen; /* to skip over non-understood msgs */
 u_char ifm_version; /* future binary compatibility */
 u_char ifm_type; /* message type */
 int ifm_addr; /* like rtm_addr */
 int ifm_flags; /* value of if_flags */
 u_short ifm_index; /* index for associated ifp */
 struct if_data ifm_dat /* statistics and other data about if */
};

struct ifa_msghdr {
 u_short ifam_msglen; /* to skip over non-understood msgs */
 u_char ifam_version; /* future binary compatibility */
 u_char ifam_type; /* message type */
 int ifam_addr; /* like rtm_addr */
 int ifam_flags; /* value of ifa_flags */
 u_short ifam_index; /* index for associated ifp */
 int ifam_metric; /* value of ifa_metric */
};
```

The RTM\_IFINFO message uses an `if_msghdr` header. The RTM\_NEWADDR and RTM\_DELADDR messages use an `ifa_msghdr` header. All other messages use the `rt_msghdr` header.

The metrics structure is:

```

struct rt_metrics {
 u_long rmx_locks; /* Kernel must leave these values alone */
 u_long rmx_mtu; /* MTU for this path */
 u_long rmx_hopcount; /* max hops expected */
 u_long rmx_expire; /* lifetime for route, e.g. redirect */
 u_long rmx_recvpipe; /* inbound delay-bandwidth product */
 u_long rmx_sendpipe; /* outbound delay-bandwidth product */
 u_long rmx_ssthresh; /* outbound gateway buffer limit */
 u_long rmx_rtt; /* estimated round trip time */
 u_long rmx_rttvar; /* estimated rtt variance */
 u_long rmx_pktsent; /* packets sent using this route */
};

```

Flags include the values:

```

#define RTF_UP 0x1 /* route usable */
#define RTF_GATEWAY 0x2 /* destination is a gateway */
#define RTF_HOST 0x4 /* host entry (net otherwise) */
#define RTF_REJECT 0x8 /* host or net unreachable */
#define RTF_DYNAMIC 0x10 /* created dynamically (by redirect) */
#define RTF_MODIFIED 0x20 /* modified dynamically (by redirect) */
#define RTF_DONE 0x40 /* message confirmed */
#define RTF_MASK 0x80 /* subnet mask present */
#define RTF_CLONING 0x100 /* generate new routes on use */
#define RTF_XRESOLVE 0x200 /* external daemon resolves name */
#define RTF_LLINFO 0x400 /* generated by ARP or ESIS */
#define RTF_STATIC 0x800 /* manually added */
#define RTF_BLACKHOLE 0x1000 /* just discard pkts (during updates) */
#define RTF_PROTO2 0x4000 /* protocol specific routing flag */
#define RTF_PROTO1 0x8000 /* protocol specific routing flag */

```

Specifiers for metric values in *rmx\_locks* and *rtm\_inits* are:

```

#define RTV_MTU 0x1 /* init or lock _mtu */
#define RTV_HOPCOUNT 0x2 /* init or lock _hopcount */
#define RTV_EXPIRE 0x4 /* init or lock _expire */
#define RTV_RPIPE 0x8 /* init or lock _recvpipe */
#define RTV_SPIPE 0x10 /* init or lock _sendpipe */
#define RTV_SSTHRESH 0x20 /* init or lock _ssthresh */
#define RTV_RTT 0x40 /* init or lock _rtt */
#define RTV_RTTVAR 0x80 /* init or lock _rttvar */

```

Specifiers for which addresses are present in the messages are:

```

#define RTA_DST 0x1 /* destination sockaddr present */
#define RTA_GATEWAY 0x2 /* gateway sockaddr present */
#define RTA_NETMASK 0x4 /* netmask sockaddr present */

```

```
#define RTA_GENMASK 0x8 /* cloning mask sockaddr present */
#define RTA_IFP 0x10 /* interface name sockaddr present */
#define RTA_IFA 0x20 /* interface addr sockaddr present */
#define RTA_AUTHOR 0x40 /* sockaddr for author of redirect */
#define RTA_BRD 0x80 /* for NEWADDR, */
/* broadcast or p-p dest addr */
```

## Examples:

Use the following code to set the default route:

```
#include <sys/socket.h>
#include <sys/uio.h>
#include <unistd.h>
#include <net/route.h>
#include <netinet/in.h>
#include <stdio.h>
#include <libgen.h>
#include <arpa/inet.h>
#include <process.h>
#include <errno.h>

struct my_rt
{
 struct rt_msghdr rt;
 struct sockaddr_in dst;
 struct sockaddr_in gate;
 struct sockaddr_in mask;
};

int main(int argc, char **argv)
{
 int s;
 struct rt_msghdr *rtm;
 struct sockaddr_in *dst, *gate, *mask;
 struct my_rt my_rt;

 if(argc < 2)
 {
 fprintf(stderr, "Usage: %s: <ip_addr_of_default_gateway>\n",
 basename(argv[0]));
 return 1;
 }

 if((s = socket(AF_ROUTE, SOCK_RAW, 0)) == -1)
 {
 perror("socket");
 }
}
```

```
 return 1;
 }

 memset(&my_rt, 0x00, sizeof(my_rt));

 rtm = &my_rt.rtm;

 dst = &my_rt.dst;
 gate = &my_rt.gate;
 mask = &my_rt.mask;

 rtm->rtm_type = RTM_ADD;
 rtm->rtm_flags = RTF_UP | RTF_GATEWAY | RTF_STATIC;
 rtm->rtm_msglen = sizeof(my_rt);
 rtm->rtm_version = RTM_VERSION;
 rtm->rtm_seq = 1234;
 rtm->rtm_addrs = RTA_DST | RTA_GATEWAY | RTA_NETMASK;
 rtm->rtm_pid = getpid();

 dst->sin_len = sizeof(*dst);
 dst->sin_family = AF_INET;

 mask->sin_len = sizeof(*mask);
 mask->sin_family = AF_INET;

 gate->sin_len = sizeof(*gate);
 gate->sin_family = AF_INET;
 inet_aton(argv[1], &gate->sin_addr);

AGAIN:
 if(write(s, rtm, rtm->rtm_msglen) < 0)
 {
 if(errno == EEXIST && rtm->rtm_type == RTM_ADD)
 {
 rtm->rtm_type = RTM_CHANGE;
 goto AGAIN;
 }

 perror("write");
 return 1;
 }
 return 0;
}
```

**See also:**

*setsockopt()*, *socket()*, *sysctl()*

`npm-tcpip.so` in the *Utilities Reference*

## ***Rrcmd()***

© 2004, QNX Software Systems Ltd.

*Execute a command on a remote host (via a SOCKS server)*

### **Synopsis:**

```
int Rrcmd(char ** ahost,
 int inport,
 const char * locuser,
 const char * remuser,
 const char * cmd,
 int * fd2p);
```

### **Arguments:**

|                |                                                                                                                                                           |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ahost</i>   | The name of the host that you want to execute the command on. If the function can find the host, it sets * <i>ahost</i> to the standard name of the host. |
| <i>inport</i>  | The well-known Internet port on the host, where the server resides.                                                                                       |
| <i>locuser</i> | The user ID on the local machine.                                                                                                                         |
| <i>remuser</i> | The user ID on the remote machine.                                                                                                                        |
| <i>cmd</i>     | The command that you want to execute.                                                                                                                     |
| <i>fd2p</i>    | See <i>rcmd()</i> .                                                                                                                                       |

### **Library:**

**libsocks**

### **Description:**

The *Rrcmd()* function is a cover function for *rcmd()* — the difference is that *Rrcmd()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

**Returns:**

A valid socket descriptor; or -1 if an error occurs and a message is printed to standard error.

**Classification:**

SOCKS

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*Raccept()*, *Rbind()*, *rcmd()*, *Rconnect()*, *Rgetsockname()*, *Rlisten()*, *Rselect()*, *SOCKSinit()*

SOCKS — A Basic Firewall

## ***rresvport()***

© 2004, QNX Software Systems Ltd.

*Obtain a socket with a privileged address*

### **Synopsis:**

```
#include <unistd.h>

int rresvport(int * port);
```

### **Arguments:**

*port* An address in the privileged port space. Privileged Internet ports are those in the range 0 to 1023. Only the superuser may bind this type of address to a socket.

### **Library:**

libsocket

### **Description:**

The *rresvport()* function returns a descriptor to a socket with an address in the privileged port space. The *ruserok()* function is used by servers to authenticate clients requesting service with *rcmd()*. All three functions are present in the same file and are used by the **rshd** server (see the *Utilities Reference*), among others.

The *rresvport()* function obtains a socket with a privileged address bound to it. This socket can be used by *rcmd()* and several other functions.

### **Returns:**

A valid, bound socket descriptor, or -1 if an error occurs (*errno* is set).

### **Errors:**

The error code EAGAIN is overloaded to mean “All network ports in use.”



## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*rcmd()*, *ruserok()*

**rshd** in the *Utilities Reference*

## ***Rselect()***

© 2004, QNX Software Systems Ltd.

*Check for descriptors that are ready for reading or writing (via a SOCKS server)*

### **Synopsis:**

```
int Rselect(int width,
 fd_set * readfds,
 fd_set * writefds,
 fd_set * exceptionfds,
 struct timeval * timeout);
```

### **Arguments:**

|                     |                                                                                                                                                                                                                                                                                      |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>width</i>        | The number of descriptors to check in the given sets. Only the descriptors from 0 through ( <i>width</i> -1) in the descriptor sets are examined. Therefore, the value of <i>width</i> must be at least as large as:<br><br>( <i>highest valued file descriptor in the sets</i> ) +1 |
| <i>readfds</i>      | NULL, or a pointer to a <b>fd_set</b> object that specifies the descriptors to check for files that are ready for reading. The function replaces the set with the file descriptors that are actually ready for reading.                                                              |
| <i>writefds</i>     | NULL, or a pointer to a <b>fd_set</b> object that specifies the descriptors to check for files that are ready for writing. The function replaces the set with the file descriptors that are actually ready for writing.                                                              |
| <i>exceptionfds</i> | NULL, or a pointer to a <b>fd_set</b> object that specifies the descriptors to check for files that have an exceptional condition pending. The function replaces the set with the file descriptors that actually have an exceptional condition pending.                              |
| <i>timeout</i>      | NULL, or a pointer to a <b>timeval</b> that specifies how long to wait for the selection to complete.                                                                                                                                                                                |

## Library:

`libsocks`

## Description:

The *Rselect()* function is a cover function for *select()* — the difference is that *Rselect()* does its job via a SOCKS server.

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

## Returns:

The number of ready descriptors in the descriptor sets, 0 if the *timeout* expired, or -1 if an error occurs (*errno* is set).

## Classification:

SOCKS

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*select()*

SOCKS — A Basic Firewall

## ***rsrcdbmgr\_attach()***

© 2004, QNX Software Systems Ltd.

*Reserve a system resource for a process*

### **Synopsis:**

```
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int rsrcdbmgr_attach(rsrc_request_t * list,
 int count);
```

### **Arguments:**

*list*      An array of `rsrc_request_t` structures that describe the resources that you want to reserve; see below.

*count*     The number of entries in the array.

### **Library:**

`libc`

### **Description:**

The resource database manager allocates and keeps track of system resources i.e. it manages these resources. The system resources currently tracked are:

- memory
- IRQs
- DMA channels
- I/O ports.

Major and minor device numbers are handled with separate `rsrcdbmgr_devno_attach()` and `rsrcdbmgr_devno_detach()` functions.

There are two main functions that drivers can use to communicate with the resource database:

- `rsrcdbmgr_attach()`
- `rsrcdbmgr_detach()`

The *rsrddbmgr\_attach()* function reserves a resource range(s) from the database of available resources for a process. Other processes can't reserve this resource range until the resource is returned to the system (usually with the *rsrddbmgr\_detach()* call). The requested resources are returned in a *list* of `_rsrc_request` structures with the *start* and *end* fields filled in. The number of resources requested is specified in *count*.




---

Reserving the resources doesn't give you access to them; you still have to use *mmap()*, *InterruptAttach()*, or another means.

---

When you're finished with the resource, you must return it to the system. The easiest way to return the resource is to call *rsrddbmgr\_detach()* with the same *start*, *end*, and type (via the *flags* field) that were issued for the resource.

### **rsrc\_request\_t structure**

The resource requests structure looks like this:

```
typedef struct _rsrc_request {
 uint64_t length;
 uint64_t align;
 uint64_t start;
 uint64_t end;
 uint32_t flags;
 uint32_t zero[3]; /* Reserved */
} rsrc_request_t;
```

The members include:

- |                   |                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>length</i>     | The length of the resource that you want to reserve. You must set this member.                                                                                                             |
| <i>align</i>      | The alignment of the resource.                                                                                                                                                             |
| <i>start, end</i> | The range of resource that you want to reserve.                                                                                                                                            |
| <i>flags</i>      | The type of the resource, as well as flags that affect the request. You must set this member to be one of the following resource types (defined in <code>&lt;sys/rsrddbmgr.h&gt;</code> ): |

- RSRCDBMGR\_DMA\_CHANNEL — DMA channel
- RSRCDBMGR\_IO\_PORT — I/O port address
- RSRCDBMGR\_IRQ — Interrupt address
- RSRCDBMGR\_MEMORY — Memory address
- RSRCDBMGR\_PCLMEMORY — PCI memory address

You can OR in the following flags (also defined in `<sys/rsrcdbmgr.h>`):

- RSRCDBMGR\_FLAG\_ALIGN — the contents of the *align* field are valid, and the requested resource starts with the given alignment.
- RSRCDBMGR\_FLAG\_RANGE — the contents of the *start* and *end* fields are valid, and the requested resource is in the range *start* to *end*, inclusive.
- RSRCDBMGR\_FLAG\_SHARE — other processes can have access to an allocated resource.
- RSRCDBMGR\_FLAG\_TOPDOWN — start the search for a free resource block from *end*. If you also set RSRCDBMGR\_FLAG\_RANGE, this flag makes the search start from the *end* of the available range.

## Returns:

EOK, or -1 if an error occurred (*errno* is set).

## Errors:

|        |                                                           |
|--------|-----------------------------------------------------------|
| EAGAIN | The resource request can't be filled.                     |
| EINVAL | Invalid argument.                                         |
| ENOMEM | Insufficient memory to allocate internal data structures. |

**Examples:**

When you start the system, the startup code and special programs that know how to probe the hardware call *rsrddbmgr\_create()* to register the hardware in the resource database. The following examples don't do this seeding, so they'll fail with an error code of EINVAL.

**Example 1**

```

/*
 * Request one DMA Channel, with length 1, from the
 * entire range of available DMA channel resources.
 */
#include <stdio.h>
#include <sys/rsrddbmgr.h>
#include <sys/rsrddbmsg.h>

int main(int argc, char **argv) {
 int count;
 rsrc_request_t req;

 memset(&req, 0, sizeof(req));
 req.length = 1;
 req.flags = RSRDCBMGR_DMA_CHANNEL;
 count = 1;

 if (rsrddbmgr_attach(&req, count) == -1) {
 perror("Problem attaching to resource ");
 exit(1);
 }

 printf("You can use DMA channel 0x%llx \n",
 req.start);

 ...
 /* Do something with the acquired resource */
 ...

 /* To return the resource to the database: */
 if (rsrddbmgr_detach(&req, count) == -1) {
 perror("Problem detaching resource \n");
 exit(1);
 }

 return(0);
}

```

## Example 2

```
/*
 * Request memory that's 4-byte aligned
 * and has a length of 50.
 */

#include <stdio.h>
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int main(int argc, char **argv) {
 int count;
 rsrc_request_t req;

 memset(&req, 0, sizeof(req));
 req.align = 4;
 req.length = 50;
 req.flags = RSRCDBMGR_FLAG_ALIGN | RSRCDBMGR_MEMORY;
 count = 1;

 if (rsrcdbmgr_attach(&req, count) == -1) {
 perror("Problem attaching to resource ");
 exit(1);
 }

 printf("You can use memory from 0x%llx 0x%llx inclusive. \n",
 req.start, req.end);

 ...
 /* Do something with the acquired resource */
 ...

 /* To return the resource to the database: */
 if (rsrcdbmgr_detach(&req, count) == -1) {
 perror("Problem detaching resource \n");
 exit(1);
 }

 return(0);
}
```

## Example 3

```
/*
 * Request two resources:
 * I/O port 0 and an IRQ in the range 10-12
 * from the available resources.
```



```
*/
#include <stdio.h>
#include <sys/rsrddbmgr.h>
#include <sys/rsrddbmsg.h>

int main(int argc, char **argv) {
 int count;
 rsrc_request_t req[2];

 memset(req, 0, 2*sizeof(*req));
 req[0].start = 0;
 req[0].end = 0;
 req[0].length = 1;
 req[0].flags = RSRDCBMGR_FLAG_RANGE | RSRDCBMGR_IO_PORT;

 req[1].start = 10;
 req[1].end = 12;
 req[1].length = 1;
 req[1].flags = RSRDCBMGR_FLAG_RANGE | RSRDCBMGR_IRQ;
 count = 2;

 if (rsrddbmgr_attach(req, count) == -1) {
 perror("Problem attaching to resource ");
 exit(1);
 }

 printf("You can use io-port 0x%llx \n",
 req[0].start);
 printf("You can use irq 0x%llx \n",
 req[1].start);

 ...
 /* Do something with the acquired resource */
 ...

 /* To return the resource to the database: */
 if (rsrddbmgr_detach(req, count) == -1) {
 perror("Problem detaching resource \n");
 exit(1);
 }

 return(0);
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*rsrcdbmgr\_create()*, *rsrcdbmgr\_detach()*, *rsrcdbmgr\_destroy()*,  
*rsrcdbmgr\_devno\_attach()*, *rsrcdbmgr\_devno\_detach()*,  
*rsrcdbmgr\_query()*

### **Synopsis:**

```
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int rsrcdbmgr_create(rsrc_alloc_t *item,
 int count);
```

### **Arguments:**

*item*      An array of `rsrc_alloc_t` structures that describe the resources that you want to create; see below.

*count*     The number of entries in the array.

### **Library:**

`libc`

### **Description:**

The `rsrcdbmgr_create()` function creates one or more system resources. If the function completes successfully, *count* resources are returned in *item*.

#### **`rsrc_alloc_t` structure**

The structure of a basic resource request looks like this:

```
typedef struct _rsrc_alloc {
 uint64_t start; /* Start of resource range */
 uint64_t end; /* End of resource range */
 uint32_t flags; /* Resource type | Resource flags */
} rsrc_alloc_t;
```

The members include:

*start, end*      The resource range.

*flags*      The type of the resource, as well as flags that affect the request. You must set this member to be one of the following resource types (defined in `<sys/rsrcdbmgr.h>`):

- RSRCDMGR\_DMA\_CHANNEL — DMA channel
- RSRCDMGR\_IO\_PORT — I/O port address
- RSRCDMGR\_IRQ or RSRMGR\_IRQ — interrupt address
- RSRCDMGR\_MEMORY — Memory address
- RSRCDMGR\_PCI\_MEMORY — PCI memory address

You can OR in the following flag (also defined in `<sys/rsrcdbmgr.h>`):

- RSRCDMGR\_FLAG\_RSVP — create and reserve a resource with a higher priority than other resources. The resource is given out only when there are no other valid ranges available.

You must set all the members.

## Returns:

EOK, or -1 if an error occurred (*errno* is set).

## Errors:

|        |                                                           |
|--------|-----------------------------------------------------------|
| EAGAIN | The resource request can't be created.                    |
| EINVAL | Invalid argument.                                         |
| ENOMEM | Insufficient memory to allocate internal data structures. |

**Examples:**

```

/*
 * Create two resources:
 * 0-4K memory allocation and 5 DMA channels.
 */
#include <stdio.h>
#include <sys/rsrddbmgr.h>
#include <sys/rsrddbmsg.h>

int main(int argc, char **argv) {
 rsrc_alloc_t alloc[2];

 memset(alloc, 0, 2* sizeof(*alloc));
 alloc[0].start = 0;
 alloc[0].end = 4*1024;
 alloc[0].flags = RSRDDBMGR_MEMORY;

 alloc[1].start = 1;
 alloc[1].end = 5;
 alloc[1].flags = RSRDDBMGR_DMA_CHANNEL;

 /* Allocate resources to the system. */
 if (rsrddbmgr_create(alloc, 2) == -1) {
 perror("Problem creating resources \n");
 exit(1);
 }

 ...
 /* Do something with the created resource */
 ...

 /* Remove the allocated resources. */
 rsrcddbmgr_destroy (alloc, 2);

 return(0);
}

```

**Classification:**

QNX Neutrino

**Safety**

Cancellation point Yes

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*rsrcdbmgr\_attach()*, *rsrcdbmgr\_destroy()*

### **Synopsis:**

```
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int rsrcdbmgr_destroy(rsrc_alloc_t *item,
 int count);
```

### **Arguments:**

*item*      An array of `rsrc_alloc_t` structures that describe the resources that you want to destroy. For more information about this structure, see the documentation for `rsrcdbmgr_create()`.

*count*     The number of entries in the array.

### **Library:**

`libc`

### **Description:**

The `rsrcdbmgr_destroy()` function removes *count* system resources that are defined in the array *item*.

### **Returns:**

EOK      Success.

-1      An error occurred; *errno* is set.

### **Errors:**

EINVAL    Invalid argument, or the resource is in use.

ENOMEM    Insufficient memory to allocate internal data structures.

**Examples:**

See the example in *rsrcdbmgr\_create()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*rsrcdbmgr\_attach()*, *rsrcdbmgr\_create()*, *rsrcdbmgr\_detach()*



### **Synopsis:**

```
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int rsrcdbmgr_detach(rsrc_request_t *list,
 int count);
```

### **Arguments:**

*list*      An array of `rsrc_request_t` structures that describe the resources that you want to return. For information about this structure, see the documentation for `rsrcdbmgr_attach()`.

*count*     The number of entries in the array.

### **Library:**

`libc`

### **Description:**

The `rsrcdbmgr_detach()` function returns *count* resources in *list* to the database of available system resources. You must return the resource with the same *start*, *end*, and *flags* (type) that were issued for the resource when it was allocated with `rsrcdbmgr_attach()`.

### **Returns:**

EOK        Success.

-1         An error occurred; *errno* is set.

### **Errors:**

EINVAL     Invalid argument, or the resource is in use by a process, isn't found in the database, or can't be returned to the system.

ENOMEM     Insufficient memory to allocate internal data structures.

**Examples:**

See the examples in *rsrcdbmgr\_attach()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*rsrcdbmgr\_attach()*, *rsrcdbmgr\_destroy()*

**Synopsis:**

```
#include <sys/rsrddbmgr.h>
#include <sys/rsrddbmsg.h>

dev_t rsrddbmgr_devno_attach(const char * name,
 int minor_request,
 int flags);
```

**Arguments:**

*name*                    The name of the class of devices that you want to get the major number for. This string can be anything, but the following class names are defined in `<sys/fstype.h>`:

| Constant                       | Value                  | Class                                                                               |
|--------------------------------|------------------------|-------------------------------------------------------------------------------------|
| <code>_MAJOR_PATHMGR</code>    | <code>"pathmgr"</code> | Used only by the path manager                                                       |
| <code>_MAJOR_DEV</code>        | <code>"dev"</code>     | Devices in <code>/dev</code> with only one instance (e.g. <code>/dev/tty</code> )   |
| <code>_MAJOR_BLK_PREFIX</code> | <code>"blk-"</code>    | All block devices (e.g. <code>/dev/hd[0-9]*</code> would be <code>"blk-hd"</code> ) |

*continued...*

| Constant                        | Value                  | Class                                                                                        |
|---------------------------------|------------------------|----------------------------------------------------------------------------------------------|
| <code>_MAJOR_CHAR_PREFIX</code> | <code>"char - "</code> | All character devices (e.g. <code>/dev/ser [0-9] *</code> would be <code>"char-ser"</code> ) |
| <code>_MAJOR_FSYS</code>        | <code>"fsys"</code>    | All filesystems                                                                              |

*minor\_request* The minor device number that you want to reserve, or -1 to let the system assign the next available minor number.

*flags* Presently, there are no flags; pass zero for this argument.

**Library:**

`libc`

**Description:**

The function *rsrddbmgr\_devno\_attach()* reserves a device number that consists of:

- a major number that corresponds to the given device class. If there isn't already a major number associated with the class, a new major number is assigned to it.
- a minor number that's based on *minor\_request*. If *minor\_request* is -1, the function returns the first free minor number in the specified class.

There's a maximum of 64 major numbers (0 through 63) on the system, and a maximum of 1024 minor numbers (0 through 1023) per major number.

Major and minor numbers are used only by resource managers and are exposed through the *rdev* member of the *iofunc\_attr\_t* structure, and correspondingly the *st\_rdev* member of the *stat* structure. They aren't required for proper operation; on simple devices, an entry will be simulated for you.

## Returns:

A *dev\_t* object that contains the major and minor numbers, or -1 if an error occurs (*errno* is set).

You can extract the major and minor number values from the *dev\_t* object by using the *major()* and *minor()* macros defined in *<sys/types.h>*. For more information, see the documentation for *stat()*.

## Errors:

EINVAL Invalid argument.

## Examples:

```
#include <sys/rsrddbmgr.h>
#include <sys/rsrddbmsg.h>

char *dev_name;
int myminor_request, flags=0;
dev_t major_minor;

major_minor = rsrddbmgr_devno_attach
 (dev_name, myminor_request, flags);

:

rsrddbmgr_devno_detach(major_minor, flags);
```

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*iofunc\_attr\_t*, *rsrdbmgr\_attach()*, *rsrdbmgr\_devno\_detach()*,  
*stat()*

### **Synopsis:**

```
#include <sys/rsrdbmgr.h>
#include <sys/rsrdbmsg.h>

int rsrdbmgr_devno_detach(dev_t devno,
 int flags);
```

### **Arguments:**

*devno*     A **dev\_t** object that was returned by *rsrdbmgr\_devno\_attach()*.

*flags*     Presently, there are no flags; pass zero for this argument.

### **Library:**

**libc**

### **Description:**

The function *rsrdbmgr\_devno\_detach()* detaches device number that was attached with *rsrdbmgr\_devno\_attach()*.

### **Returns:**

EOK        Success.

-1         An error occurred.

### **Examples:**

```
#include <sys/rsrdbmgr.h>
#include <sys/rsrdbmsg.h>

dev_t dev_num;
int flags=0;

:

rsrdbmgr_devno_detach(dev_num, flags);
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*rsrdbmgr\_attach()*, *rsrdbmgr\_devno\_attach()*



### **Synopsis:**

```
#include <sys/rsrcdbmgr.h>
#include <sys/rsrcdbmsg.h>

int rsrcdbmgr_query(rsrc_alloc_t *list,
 int listcnt,
 int start,
 uint32_t type);
```

### **Arguments:**

- item* NULL, or an array of `rsrc_alloc_t` structures that the function can fill with information about the resources that it finds. For more information about this structure, see the documentation for `rsrcdbmgr_create()`.
- listcnt* The number of entries in the array.
- start* The index that you want to start searching at.
- type* The type of resource that you want to query; one of the following (defined in `<sys/rsrcdbmgr.h>`):
- `RSRCDBMGR_DMA_CHANNEL` — DMA channel
  - `RSRCDBMGR_IO_PORT` — I/O port address
  - `RSRCDBMGR_IRQ` or `RSRCMGR_IRQ` — interrupt address
  - `RSRCDBMGR_MEMORY` — Memory address
  - `RSRCDBMGR_PCI_MEMORY` — PCI memory address

### **Library:**

`libc`

## Description:

The *rsrcdbmgr\_query()* function queries the database for *listcnt* count of *type* resources in use, beginning at the index *start*. If you make the query with a non-NULL *list*, then the function stores a maximum of found *listcnt* resources in the array.

## Returns:

If *list* is NULL or *listcnt* is 0, then the function returns the number of resources of *type* in the database.

If *list* is non-NULL, then the function returns the number of *type* resources available in the system.

If an error occurs, the function returns -1 and sets *errno*.

## Errors:

- EINVAL      Invalid argument
- ENOMEM     Insufficient memory to allocate internal data structures.

## Examples:

List all of the memory resource blocks available in the system:

```
rsrc_alloc_t list[20];
int size, count = 0, start = 0;

while (1) {
 count = rsrcdbmgr_query(&list, 20, start, RSRCDMGR_MEMORY);
 if (count == -1)
 break;

 size = min(count-start, 20); /* In case more than 20 blocks
 were returned. */
 printf("Retrieved %d of a possible %d resource blocks", \
 size, count);

 for (count=0; count<size; count++) {
 printf("RSRC[%d] Start %d End %d \n", \
 start+count, list[count].start, list[count].end);
 }
 start += size; /* Loop again, in case there are more than
```

```
 20 blocks. */
 }
```

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*rsrddbmgr\_attach()*

## ***ruserok()***

© 2004, QNX Software Systems Ltd.

*Check the identity of a remote host*

### **Synopsis:**

```
#include <unistd.h>

int ruserok(char * rhost,
 int superuser,
 char * ruser,
 char * luser);
```

### **Arguments:**

|                  |                                                                      |
|------------------|----------------------------------------------------------------------|
| <i>rhost</i>     | The name of the remote host, as returned by <i>gethostbyaddr()</i> . |
| <i>superuser</i> | Nonzero if the local user is the superuser, zero otherwise.          |
| <i>ruser</i>     | The name of the remote user.                                         |
| <i>luser</i>     | The name of the local user.                                          |

### **Library:**

**libsocket**

### **Description:**

The *ruserok()* routine checks the identity of a remote host. It's used by servers to authenticate clients requesting service with *rcmd()*.

The *rcmd()*, *rresvport()*, and *ruserok()* functions are used by the **rshd** server (see the *Utilities Reference*), among others.

The *ruserok()* function takes a remote host's name (as returned by the *gethostbyaddr()* routine), two user names, and a flag indicating whether the local user's name is that of the superuser. Then, if the user is *not* the superuser, it checks the file **/etc/hosts.equiv** (described in the *Utilities Reference*).

If that lookup isn't done, or is unsuccessful, the **.rhosts** file in the local user's home directory is checked to see if the request for service

is allowed. If this file is owned by anyone other than the user or the superuser, or if it's writable by anyone other than the owner, the check automatically fails.

If the local domain obtained from *gethostname()* is the same as the remote domain, only the machine name need be specified.

### Returns:

- 0      The machine name is listed in the `hosts.equiv` file, or the host and remote username were found in the `.rhosts` file.
- 1     An error occurred (*errno* is set).

### Errors:

The error code EAGAIN is overloaded to mean "All network ports in use."

### Classification:

Unix

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

### See also:

*gethostbyaddr()*, *gethostname()*, *rcmd()*, *rresvport()*  
*/etc/hosts.equiv*, *rshd* in the *Utilities Reference*

## **sbrk()**

© 2004, QNX Software Systems Ltd.

*Set the allocation break value*

---

### **Synopsis:**

```
#include <unistd.h>

void* sbrk(int increment);
```

### **Arguments:**

*increment*      The amount by which to increase the current break value. This increment may be positive or negative.

### **Library:**

libc

### **Description:**

The *break* value is the address of the first byte of unallocated memory. When a program starts execution, the break value is placed following the code and constant data for the program. As memory is allocated, this pointer advances when there is no free block large enough to satisfy an allocation request. The *sbrk()* function sets a new break value for the program by adding the value of *increment* to the current break value.

The variable *\_amblksiz* (defined in *<stdlib.h>*) contains the default increment. This value may be changed by a program at any time.

### **Returns:**

A pointer to the start of the new block of memory for success, or *-1* if an error occurs (*errno* is set).

### **Errors:**

EAGAIN      The total amount of system memory available for allocation to this process is temporarily insufficient. This may occur although the space requested is less than the maximum data segment size.

**ENOMEM** The requested change allocated more space than allowed, is impossible since there's insufficient swap space available, or it caused a memory allocation conflict.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

#define alloc(x, y) y = sbrk(x);

int main(void)
{
 void* brk;

 brk = sbrk(0x3100);
 printf("New break value after sbrk(0x3100) \t%p\n",
 brk);

 brk = sbrk(0x0200);
 printf("New break value after sbrk(0x0200) \t%p\n",
 brk);

 return EXIT_SUCCESS;
}
```

## Classification:

Legacy Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*\_amblksiz, \_btext, \_edata, \_end, \_etext, brk(), calloc(), errno, free(), malloc(), realloc()*



## Synopsis:

```
#include <math.h>

double scalb(double x,
 double n);

float scalbf(float x,
 float n);
```

## Arguments:

- x* The floating point number that you want to multiply by the exponent.
- n* The exponent to apply to the radix of the machine's floating-point arithmetic.

## Library:

libm

## Description:



---

We recommend that you use *scalbn()* since it computes by exponent manipulation rather than mock multiplications or additions.

---

These functions compute  $x \times r^n$ , where  $r$  is the radix of the machine's floating point arithmetic and  $n$  is a finite number. When  $r$  is 2, *scalb()* is equivalent to *ldexp()*.

## Returns:

$x \times r^n$



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

## Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b, c, d;

 a = 10;
 b = 2;
 c = scalb(a, b);
 d = sqrt(c/a);
 printf("Radix of machines fp arithmetic is %f \n", d);
 printf("So %f = %f * (%f ^ %f) \n", c, a, d, b);

 return(0);
}
```

produces the output:

```
Radix of machines fp arithmetic is 2.000000
So 40.000000 = 10.000000 * (2.000000 ^ 2.000000)
```

## Classification:

*scalb()* is standard Unix; *scalbf()* is ANSI (draft)

### Safety

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*ldexp()*, *scalbn()*

## ***scalbn()*, *scalbnf()***

© 2004, QNX Software Systems Ltd.

*Load the exponent of a radix-independent floating point number*

### **Synopsis:**

```
#include <math.h>

double scalbn (double x,
 int n);

float scalbnf (float x,
 int n);
```

### **Arguments:**

- x*     The floating point number that you want to multiply by the exponent.
- n*     The exponent to apply to the radix of the machine's floating-point arithmetic.

### **Library:**

libm

### **Description:**

The *scalbn()* and *scalbnf()* functions compute  $x \times r^n$ , where *r* is the radix of the machine's floating point arithmetic.

### **Returns:**

$x \times r^n$



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

**Examples:**

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b, c, d;

 a = 10;
 b = 2;
 c = scalbn(a, b);
 d = sqrt(c/a);
 printf("Radix of machines fp arithmetic is %f \n", d);
 printf("So %f = %f * (%f ^ %f) \n", c, a, d, b);

 return(0);
}
```

produces the output:

```
Radix of machines fp arithmetic is 2.000000
So 40.000000 = 10.000000 * (2.000000 ^ 2.000000)
```

**Classification:**

*scalbn()* is standard Unix; *scalbnf()* is ANSI (draft)

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*scalb()*

## Synopsis:

```
#include <malloc.h>

void* _scalloc(size_t size);
```

## Arguments:

*size*     The number of bytes to allocate.

## Library:

`libc`

## Description:

The `_scalloc()` functions allocate space for an array of length *size* bytes. Each element is initialized to 0.

You must use `_sfree()` to deallocate the memory allocated by `_scalloc()`.

## Returns:

A pointer to the start of the allocated memory, or NULL if there's insufficient memory available or if the *size* is zero.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*calloc(), free(), realloc(), \_sfree(), \_smalloc(), \_srealloc()*



**Synopsis:**

```
#include <sys/types.h>
#include <sys/dir.h>

int scandir(char * dirname,
 struct direct * (* namelist []),
 int (*select)(struct dirent *),
 int (*compar)(const void *,const void *));
```

**Arguments:**

- dirname*     The name of the directory that you want to scan.
- namelist*    A pointer to a location where *scandir()* can store a pointer to the array of directory entries that it builds.
- select*      A pointer to a user-supplied subroutine that *scandir()* calls to select which entries to included in the array. The *select* routine is passed a pointer to a directory entry and should return a nonzero value if the directory entry is to be included in the array.
- If *select* is NULL, all the directory entries are included.
- compar*      A pointer to a user-supplied subroutine that's passed to *qsort()* to sort the completed array. If this pointer is NULL, the array isn't sorted.
- You can use *alphasort()* as the *compar* parameter to sort the array alphabetically.

**Library:**

`libc`

**Description:**

The *scandir()* function reads the directory *dirname* and builds an array of pointers to directory entries, using *malloc()* to allocate the space.

The *scandir()* function returns the number of entries in the array, and stores a pointer to the array in the location referenced by *namelist*.

You can deallocate the memory allocated for the array by calling *free()*. Free each pointer in the array, and then free the array itself.

## Returns:

The number of entries in the array, or -1 if the directory can't be opened for reading, or *malloc()* can't allocate enough memory to hold all the data structures.

## Classification:

Legacy Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*alphasort()*, *closedir()*, *free()*, *malloc()*, *opendir()*, *qsort()*, *readdir()*, *rewinddir()*, *seekdir()*, *telldir()*

## Synopsis:

```
#include <stdio.h>
```

```
int scanf(const char* format,
 ...);
```

## Arguments:

*format* A string that controls the format of the input, as described below. The formatting string determines what additional arguments you need to provide.

## Library:

libc

## Description:

The *scanf()* function scans input from *stdin* under control of the *format* argument, assigning values to the remaining arguments.

### Format control string

The format control string consists of zero or more *format directives* that specify what you consider to be acceptable input data. Subsequent arguments are pointers to various types of objects that the function assigns values to as it processes the format string.

A format directive can be a sequence of one or more whitespace characters or:

#### *multibyte characters*

Any character in the *format* string, other than a whitespace character or the percent character (%), that isn't part of a conversion specifier.

#### *conversion specifiers*

A sequence of characters in the *format* string that begins with a percent character (%) and is followed by:

- an optional *assignment suppression indicator*: the asterisk character (\*)
- an optional decimal integer that specifies the *maximum field width* to be scanned for the conversion
- an optional *type length* specification; one of **h**, **L**, or **l**
- a character that specifies the type of conversion to be performed; one of the characters: **c**, **d**, **e**, **f**, **g**, **i**, **n**, **o**, **p**, **s**, **u**, **X**, **x**, [

As each format directive in the format string is processed, the directive may successfully complete, fail because of a lack of input data, or fail because of a matching error as defined by the directive.

If end-of-file is encountered on the input data before any characters that match the current directive have been processed (other than leading whitespace, where permitted), the directive fails for lack of data.

If end-of-file occurs after a matching character has been processed, the directive is completed (unless a matching error occurs), and the function returns without processing the next directive.

If a directive fails because of an input character mismatch, the character is left unread in the input stream.

Trailing whitespace characters, including newline characters, aren't read unless matched by a directive. When a format directive fails, or the end of the format string is encountered, the scanning is completed, and the function returns.

When one or more whitespace characters (space, horizontal tab `\t`, vertical tab `\v`, form feed `\f`, carriage return `\r`, newline or linefeed `\n`) occur in the format string, input data up to the first non-whitespace character is read, or until no more data remains. If no whitespace characters are found in the input data, the scanning is complete, and the function returns.

An ordinary character in the format string is expected to match the same character in the input stream.

## Conversion specifiers

A conversion specifier in the format string is processed as follows:

- For conversion types other than **l**, **c** and **n**, leading whitespace characters are skipped.
- For conversion types other than **n**, all input characters, up to any specified maximum field length, that can be matched by the conversion type are read and converted to the appropriate type of value; the character immediately following the last character to be matched is left unread; if no characters are matched, the format directive fails.
- Unless you specify the assignment suppression indicator (**\***), the result of the conversion is assigned to the object pointed to by the next unused argument (if assignment suppression was specified, no argument is skipped); the arguments must correspond in number, type and order to the conversion specifiers in the format string.

## Type length specifiers

A type length specifier affects the conversion as follows:

- **hh** causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** format conversion to assign the converted value to an object of type **signed char** or **unsigned char**.
- **h** causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** (integer) format conversion to assign the converted value to an object of type **short** or **unsigned short**.
- **j** causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** conversion to assign the converted value to an object of type **intmax\_t** or **uintmax\_t**.
- **l** (“**l**”) causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** (integer) conversion to assign the converted value to an object of type **long** or **unsigned long**.
- **l** (“**l**”) causes an **a**, **A**, **e**, **E**, **f**, **F**, **g** or **G** conversion to assign the converted value to an object of type **double**.

- **l** (“el”) causes a **c**, **s** or **[** conversion to assign the converted value to an object of type **wchar\_t**.
- **ll** (double “el”) causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** format conversion to assign the converted value to an object of type **long long** or **unsigned long long**.
- **L** causes an **a**, **A**, **e**, **E**, **f**, **F**, **g** or **G** conversion to assign the converted value to an object of type **long double**.
- **t** causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** conversion to assign the converted value to an object of type **ptrdiff\_t** or to the corresponding **unsigned** type.
- **z** causes a **d**, **i**, **o**, **u**, **x**, **X** or **n** conversion to assign the converted value to an object of type **size\_t** or to the corresponding signed integer type.

### Conversion type specifiers

The valid conversion type specifiers are:

**a**, **A**, **e**, **E**, **f**, **F**, **g** or **G**

A floating-point number, infinity, or NaN, all of which have a format as expected by *strtod()*. The argument is assumed to point to an object of type **float**.

**c** Any sequence of characters in the input stream of the length specified by the field width, or a single character if you don't specify a field width. The argument is assumed to point to the first element of a character array of sufficient size to contain the sequence, *without* a terminating NUL character (`'\0'`). For a single character assignment, a pointer to a single object of type **char** is sufficient.

When an **l** (“el”) qualifier is present, a sequence of characters are converted from the initial shift state to **wchar\_t** wide characters as if by a call to *mbrtowc()*. The conversion state is described by a **mbstate\_t** object.

- d** A decimal integer with a format as expected by *strtol()* and a *base* of 10. The argument is assumed to point to an object of type **int**.
- i** An optionally signed integer with a format as expected by *strtol()* and a *base* of 0. The argument is assumed to point to an object of type **int**.
- n** No input data is processed. Instead, the number of characters that have already been read is assigned to the object of type **int** that's pointed to by the argument. The number of items that have been scanned and assigned (the return value) isn't affected by the **n** conversion type specifier.
- o** An optionally signed octal integer with a format as expected by *strtoul()* and a *base* of 8. The argument is assumed to point to an object of type **int**.
- p** A hexadecimal integer, as described for **x** conversions below. The converted value is taken as a **void \*** and then assigned to the object pointed to by the argument.
- s** A sequence of non-whitespace characters. The argument is assumed to point to the first element of a character array of sufficient size to contain the sequence of **char**, **signed char** or **unsigned char** and a terminating NUL character, which by the conversion operation adds.  
  
When an **l** ("el") qualifier is present, a sequence of characters are converted from the initial shift state to **wchar\_t** wide characters as if by a call to *mbrtowc()*. The conversion state is described by a **mbstate\_t** object.
- u** An unsigned decimal integer, consisting of one or more decimal digits. The argument is assumed to point to an object of type **unsigned int**.
- x, X** A hexadecimal integer, with a format as expected by *strtoul()* when *base* is 16. The argument is assumed to point to an object of type **unsigned**.

- [ Matches the *scanset*, a nonempty sequence of characters. The argument is assumed to point to the first element of a character array of sufficient size to contain the sequence and a terminating NUL character, which by the conversion operation adds.  
  
When an **l** (“*el*”) qualifier is present, a sequence of characters are converted from the initial shift state to **wchar\_t** wide characters as if by a call to *mbrtowc()* with *mbstate* set to **0**. The argument is assumed to point to the first element of a **wchar\_t** array of sufficient size to contain the sequence and a terminating NUL character, which the conversion operation adds.  
  
The conversion specification includes all characters in the scanlist between the beginning [ and the terminating ]. If the conversion specification starts with [^, the scanlist matches all the characters that *aren't* in the scanlist. If the conversion specification starts with [] or [^], the ] is included in the scanlist. (To scan for ] only, specify %[].)  
  
% A % character (The entire specification is %%).

A conversion type specifier of % is treated as a single ordinary character that matches a single % character in the input data. A conversion type specifier other than those listed above causes scanning to terminate, and the function to return with an error.

## Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning stopped by reaching the end of the input stream before storing any values.

## Examples:

The line:

```
scanf("%s%f%3hx%d", name, &hexnum, &decnum)
```



with input:

```
some_string 34.555e-3 abc1234
```

copies "**some\_string**" into the array *name*, skips **34.555e-3**, assigns **0xabc** to *hexnum* and **1234** to *decnum*. The return value is 3.

The program:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 char string1[80], string2[80];

 memset(string1, 0, 80);
 memset(string2, 0, 80);

 scanf("[%abcdefghijklmnopqrstuvwxy"
 "ABCDEFGHIJKLMNPOQRSTUVWXYZ]%*2s%[\n]",
 string1, string2);

 printf("%s\n", string1);
 printf("%s\n", string2);

 return EXIT_SUCCESS;
}
```

with input:

```
They may look alike, but they don't perform alike.
```

assigns "**They may look alike**" to *string1*, skips the comma (the "**%\*2s**" matches only the comma; the following blank terminates that field), and assigns " **but they don't perform alike.**" to *string2*.

To scan a date in the form "Friday March 26 1999":

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
```

```
int main(void)
{
 int day, year;
 char weekday[10], month[12];
 int retval;

 memset(weekday, 0, 10);
 memset(month, 0, 12);

 retval = scanf("%s %s %d %d",
 weekday, month, &day, &year);
 if(retval != 4) {
 printf("Error reading date.\n");
 printf("Format is: Friday March 26 1999\n");

 return EXIT_FAILURE;
 }

 printf("weekday: %s\n", weekday);
 printf("month: %s\n", month);
 printf("day: %d\n", day);
 printf("year: %d\n", year);

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fscanf()*, *fwscanf()* *sscanf()*, *swscanf()*, *vfscanf()*, *vwscanf()*, *vscanf()*,  
*vsscanf()*, *vswwscanf()*, *vwwscanf()*, *wscanf()*

## ***sched\_getparam()***

© 2004, QNX Software Systems Ltd.

*Get the current priority of a process*

### **Synopsis:**

```
#include <sched.h>

int sched_getparam(pid_t pid,
 struct sched_param *param);
```

### **Arguments:**

*pid*            The ID of the process whose priority you want to get, or 0 to get it for the current process.

*param*         A pointer to a **sched\_param** that the function fills with the scheduling parameters.

### **Library:**

libc

### **Description:**

The *sched\_getparam()* function gets the current priority of the process specified by *pid*, and puts it in the *sched\_priority* member of the **sched\_param** structure pointed to by *param*.

If *pid* is zero, the priority of the calling process is returned.

### **Returns:**

0            Success

-1          An error occurred (*errno* is set).

### **Errors:**

EPERM        The calling process doesn't have sufficient privilege to get the priority.

ESRCH        The process *pid* doesn't exist.

**Examples:**

```
#include <sched.h>
#include <stdio.h>

#define PRIORITY_ADJUSTMENT 5

int main (void)
{
 int max_priority;
 struct sched_param param;

 /* Find out the MAX priority for the FIFO Scheduler */
 max_priority = sched_get_priority_max(SCHED_FIFO);

 /* Find out what the current priority is. */
 sched_getparam(0, ¶m);

 printf("The assigned priority is %d.\n", param.sched_priority);
 printf("The current priority is %d.\n", param.sched_curpriority);

 param.sched_priority = ((param.sched_curpriority +
 PRIORITY_ADJUSTMENT) <= max_priority) ?
 (param.sched_curpriority + PRIORITY_ADJUSTMENT) : -1;

 if (param.sched_priority == -1)
 {
 printf("Cannot increase the priority by %d. Try a lesser value\n",
 PRIORITY_ADJUSTMENT);
 return(0);
 }

 sched_setscheduler (0, SCHED_FIFO, ¶m);

 sched_getparam(0, ¶m);
 printf("The newly assigned priority is %d.\n", param.sched_priority);
 printf("The current priority is %d.\n", param.sched_curpriority);

 return(0);
}
```

**Classification:**

POSIX 1003.1 (Realtime Extensions)

## Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Currently, the implementation of *sched\_getparam()* isn't 100% POSIX 1003.1-1996. The *sched\_getparam()* function returns the scheduling parameters for thread 1 in the process *pid*, or for the *calling thread* if *pid* is 0.

If you depend on this in new code, *it will not be portable*. POSIX 1003.1 says *sched\_getparam()* should return -1 and set *errno* to EPERM in a multithreaded application.

## See also:

*errno*, *getprio()*, *sched\_get\_priority\_max()*, *sched\_get\_priority\_min()*, *sched\_getscheduler()*, **sched\_param**, *sched\_setparam()*, *sched\_setscheduler()*, *sched\_yield()*, *setprio()*

## ***sched\_get\_priority\_adjust()***

*Calculate the allowable priority for the scheduling policy*

### **Synopsis:**

```
#include <sched.h>

int sched_get_priority_adjust(int prio,
 int policy,
 int adjust);
```

### **Arguments:**

- prio*      The original priority value. If negative, the priority of the calling thread is used.
- policy*    The scheduling algorithm being used. The valid arguments are listed in *sched\_get\_priority\_max()*. If *policy* is SCHED\_NOCHANGE, the function uses the algorithm of the calling thread.
- adjust*    The priority change, relative to *prio*. A value of +10 results in a final priority of *prio*+10, provided that this amount of adjustment is allowed.

### **Library:**

`libc`

### **Description:**

The *sched\_get\_priority\_adjust()* function calculates the requested priority change relative to another thread and returns the allowable value.

This function makes it easier for you to set relative priorities in order to ensure proper precedence.

### **Returns:**

- >0      The allowed priority value. The value will never exceed the range of values allowed by *sched\_get\_priority\_min()* and *sched\_get\_priority\_max()*.

<0 Failure; the negative of the *errno* value.

## Errors:

EINVAL The value of the *policy* parameter doesn't represent a defined scheduling policy.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *sched\_getparam()*, *sched\_get\_priority\_max()*,  
*sched\_get\_priority\_min()*, *sched\_setparam()*, *sched\_getscheduler()*,  
*sched\_setscheduler()*



## ***sched\_get\_priority\_max()***

*Get the maximum priority for the scheduling policy*

### **Synopsis:**

```
#include <sched.h>

int sched_get_priority_max(int policy);
```

### **Arguments:**

- policy*     The scheduling policy, which must be one of:
- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority ready thread runs until it blocks or is preempted by a higher priority thread.
  - SCHED\_RR — similar to SCHED\_FIFO, except that threads at the same priority level timeslice (round robin) every 50 msec.
  - SCHED\_OTHER — currently the same as SCHED\_RR.
  - SCHED\_SPORADIC — sporadic scheduling. For more information, see *pthread\_attr\_setschedpolicy()*, and “Scheduling algorithms” in the chapter on the Neutrino microkernel in the *System Architecture* guide.

### **Library:**

`libc`

### **Description:**

The *sched\_get\_priority\_max()* function returns the maximum value for the scheduling policy specified by *policy*.

### **Returns:**

The appropriate minimum for success, or -1 (*errno* is set).

## Errors:

|        |                                                                                         |
|--------|-----------------------------------------------------------------------------------------|
| EINVAL | The value of the <i>policy</i> parameter doesn't represent a defined scheduling policy. |
| ENOSYS | The <i>sched_get_priority_max()</i> function isn't currently supported.                 |

## Examples:

See *sched\_getparam()*.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*sched\_getparam()*, *sched\_get\_priority\_min()*, *sched\_setparam()*,  
*sched\_getscheduler()*, *sched\_setscheduler()*

## ***sched\_get\_priority\_min()***

*Get the minimum priority for the scheduling policy*

### **Synopsis:**

```
#include <sched.h>

int sched_get_priority_min(int policy);
```

### **Arguments:**

- policy* The scheduling policy, which must be one of:
- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority ready thread runs until it blocks or is preempted by a higher priority thread.
  - SCHED\_RR — similar to SCHED\_FIFO, except that threads at the same priority level timeslice (round robin) every 50 msec.
  - SCHED\_OTHER — currently the same as SCHED\_RR.
  - SCHED\_SPORADIC — sporadic scheduling. For more information, see *pthread\_attr\_setschedpolicy()*, and “Scheduling algorithms” in the chapter on the Neutrino microkernel in the *System Architecture* guide.

### **Library:**

`libc`

### **Description:**

The *sched\_get\_priority\_min()* function returns the minimum value for the scheduling policy specified by *policy*.

### **Returns:**

The appropriate minimum for success, or -1 (*errno* is set).

## Errors:

|        |                                                                                         |
|--------|-----------------------------------------------------------------------------------------|
| EINVAL | The value of the <i>policy</i> parameter doesn't represent a defined scheduling policy. |
| ENOSYS | The <i>sched_get_priority_min()</i> function isn't currently supported.                 |

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_setparam()*,  
*sched\_getscheduler()*, *sched\_setscheduler()*

**Synopsis:**

```
#include <sched.h>

int sched_getscheduler(pid_t pid);
```

**Arguments:**

*pid*      The ID of the process whose scheduling policy you want to find, or zero if you want to get the policy for the current process.

**Library:**

libc

**Description:**

The *sched\_getscheduler()* function gets the current scheduling policy of process *pid*. If *pid* is zero, the scheduling policy of the calling process is returned.

**Returns:**

The scheduling policy, or -1 if an error occurred (*errno* is set).

**Errors:**

ESRCH      The process *pid* doesn't exist.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

## **Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

## **Caveats:**

Currently, the implementation of *sched\_getscheduler()* isn't 100% POSIX 1003.1-1996. The *sched\_getscheduler()* function returns the scheduling policy for thread 1 in the process *pid*, or for the *calling thread* if *pid* is 0.

If you depend on this in new code, *it will not be portable*. POSIX 1003.1 says *sched\_getscheduler()* should return -1 and set *errno* to EPERM in a multithreaded application.

## **See also:**

*errno*, *getprio()*, *sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_get\_priority\_min()*, *sched\_setparam()*, *sched\_setscheduler()*, *sched\_yield()*, *setprio()*

**Synopsis:**

```

#include <sched.h>

struct sched_param {
 int32_t sched_priority;
 int32_t sched_curpriority;
 union {
 int32_t reserved[8];
 struct {
 int32_t __ss_low_priority;
 int32_t __ss_max_repl;
 struct timespec __ss_repl_period;
 struct timespec __ss_init_budget;
 }
 __ss;
 }
 __ss_un;
}

#define sched_ss_low_priority __ss_un.__ss.__ss_low_priority
#define sched_ss_max_repl __ss_un.__ss.__ss_max_repl
#define sched_ss_repl_period __ss_un.__ss.__ss_repl_period
#define sched_ss_init_budget __ss_un.__ss.__ss_init_budget

```

**Description:**

You'll use the **sched\_param** structure when you get or set the scheduling parameters for a thread or process.

You can use these functions to get the scheduling parameters:

- *pthread\_attr\_getschedparam()*
- *pthread\_getschedparam()*
- *sched\_getparam()*
- *SchedGet()*

You can use these functions to set the scheduling parameters:

- *pthread\_attr\_setschedparam()*
- *pthread\_setschedparam()*

- *sched\_setparam()*
- *sched\_setscheduler()*
- *SchedSet()*
- *ThreadCreate()*

The members of `sched_param` include:

*sched\_priority* When you get the scheduling parameters, this member reflects the priority that was assigned to the thread or process. It doesn't reflect any temporary adjustments due to priority inheritance. When you set the scheduling parameters, set this member to the priority that you want to use. The priority must be between the minimum and maximum values returned by *sched\_get\_priority\_min()* and *sched\_get\_priority\_max()* for the scheduling policy.

*sched\_curpriority* When you get the scheduling parameters, this member is set to the priority that the thread or process is currently running at. This is the value that the kernel uses when making scheduling decisions. When you set the scheduling parameters, this member is ignored.

The other members are used with sporadic scheduling. The following `#define` directives create the POSIX names that correspond to those members and should be used instead of accessing members directly.

*sched\_ss\_low\_priority*

The background or low priority for the thread that's executing.



*sched\_ss\_max\_repl*

The maximum number of times a replenishment will be scheduled, typically because of a blocking operation. After a thread has blocked this many times, it automatically drops to the low-priority level for the remainder of its execution until its execution budget is replenished.

*sched\_ss\_repl\_period*

The time that should be used for scheduling the replenishment of an execution budget after being blocked or having overrun the maximum number of replenishments. This time is used as an offset against the time that a thread is made READY.

*sched\_ss\_init\_budget*

The time that should be used for the thread's execution budget. As the thread runs at its high-priority level, its execution time is carved out of this budget. Once the budget is entirely depleted, the thread drops to its low-priority level, where, if possible because of priority arrangements, it can continue to run until the execution budget is replenished.



- The *sched\_priority* must always be higher than *sched\_ss\_low\_priority*.
  - The *sched\_ss\_max\_repl* must be smaller than SS\_REPL\_MAX.
  - The *sched\_ss\_init\_budget* must be larger than *sched\_ss\_repl\_period*.
- 

For more information, see “Scheduling algorithms” in the Neutrino Microkernel chapter of the *System Architecture* guide.

**Examples:**

This code shows a duty-cycle usage of the sporadic server thread:

```
#include <stdio.h>
#include <errno.h>
#include <sched.h>
```

```
#include <pthread.h>
#include <inttypes.h>
#include <sys/syspage.h>
#include <sys/neutrino.h>

/* 50 % duty cycle of 5 secs on 5 secs off */
struct timespec g_init_budget = { 5, 0 };
struct timespec g_repl_period = { 10, 0 };

#define MY_HIGH_PRIORITY 5
#define MY_LOW_PRIORITY 4
#define MY_REPL_PERIOD g_repl_period
#define MY_INIT_BUDGET g_init_budget
#define MY_MAX_REPL 10

#define DUTY_CYCLE_LOOPS 10

/*
 * Run a compute bound thread (minimal blocking) to show the duty cycle.
 */
void *st_duty_check(void *arg) {
 struct sched_param params;
 uint64_t stime, etime, cps;
 double secs;
 int ret, prio;
 int prevprio, iterations;

 stime = ClockCycles();
 cps = SYSPAGE_ENTRY(qtime)->cycles_per_sec;
 iterations = 0;

 printf("\n");

 prevprio = -1;
 while(iterations < DUTY_CYCLE_LOOPS) {
 etime = ClockCycles();
 ret = pthread_getschedparam(pthread_self(), &prio, ¶ms);

 if(ret != 0) {
 printf("pthread_getschedparam() failed %d \n", errno);
 break;
 } else if (prevprio != -1 && prevprio != params.sched_priority) {
 stime = etime - stime;
 secs = (double)stime / (double)cps;
 printf("pri %d (cur %d) %lld cycles %g secs\n",
 params.sched_priority,
 params.sched_curpriority,
 stime, secs);
 stime = etime;
 iterations++;
 }
 }
}
```

```
 }
 prevprio = params.sched_priority;
}

return NULL;
}

int main(int argc, char **argv) {
 struct sched_param params;
 pthread_attr_t attr;
 pthread_t thr;
 int ret;

 /* Set the attribute structure with the sporadic values */
 printf("# Set sporadic attributes ...");
 pthread_attr_init(&attr);
 ret = pthread_attr_setinheritsched(&attr, PTHREAD_EXPLICIT_SCHED);
 if(ret != 0) {
 printf("pthread_attr_setinheritsched() failed %d \n", errno);
 return 1;
 }

 ret = pthread_attr_setschedpolicy(&attr, SCHED_SPORADIC);
 if(ret != 0) {
 printf("pthread_attr_setschedpolicy() failed %d %d\n", ret, errno);
 return 1;
 }

 params.sched_priority = MY_HIGH_PRIORITY;
 params.sched_ss_low_priority = MY_LOW_PRIORITY;
 memcpy(¶ms.sched_ss_init_budget, &MY_INIT_BUDGET, sizeof(MY_INIT_BUDGET));
 memcpy(¶ms.sched_ss_repl_period, &MY_REPL_PERIOD, sizeof(MY_REPL_PERIOD));
 params.sched_ss_max_repl = MY_MAX_REPL;
 ret = pthread_attr_setschedparam(&attr, ¶ms);
 if(ret != 0) {
 printf("pthread_attr_setschedparam() failed %d \n", errno);
 return 1;
 }
 printf("OK\n");

 /* Create a sporadic thread to check the duty cycle */
 printf("# Creating thread duty cycle thread (%d changes) ... ", DUTY_CYCLE_LOOPS);
 ret = pthread_create(&thr, &attr, st_duty_check, NULL);
 if(ret != 0) {
 printf("pthread_create() failed %d \n", errno);
 return 1;
 }
 pthread_join(thr, NULL);
 printf("OK\n");
}
```

```
return 0;
}
```

See also *sched\_getparam()*.

## Classification:

POSIX 1003.1 (Realtime Extensions)

## See also:

*pthread\_attr\_getschedparam()*, *pthread\_attr\_setschedparam()*,  
*pthread\_getschedparam()*, *pthread\_setschedparam()*,  
*sched\_getparam()*, *sched\_setparam()*, *sched\_setscheduler()*,  
*SchedGet()*, *SchedSet()*, *ThreadCreate()*

“Scheduling algorithms” in the Neutrino Microkernel chapter of the  
*System Architecture* guide

## ***sched\_rr\_get\_interval()***

*Get the execution time limit of a process*

### **Synopsis:**

```
#include <sched.h>

int sched_rr_get_interval(
 pid_t pid,
 struct timespec * interval);
```

### **Arguments:**

*pid*            The process ID whose execution time limit you want to get.

*interval*       A pointer to a **timespec** structure that the function updates with the process's current execution time limit.

### **Library:**

**libc**

### **Description:**

The *sched\_rr\_get\_interval()* function updates *interval* with the current execution time limit for the process, *pid*. If *pid* is 0, the current execution time limit for the calling process is returned.

### **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

### **Errors:**

ENOSYS      The *sched\_rr\_get\_interval()* function isn't currently supported.

ESRCH       The process *pid* can't be found.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`timespec`

## Synopsis:

```
#include <sched.h>

int sched_setparam(
 pid_t pid,
 const struct sched_param *param);
```

## Arguments:

- pid*            The ID of the process whose priority you want to set, or 0 to set it for the current process.
- param*          A pointer to a **sched\_param** structure whose *sched\_priority* member holds the priority that you want to assign to the process.

## Library:

**libc**

## Description:

The *sched\_setparam()* function changes the priority of process *pid* to that of the *sched\_priority* member in the **sched\_param** structure pointed to by *param*. If *pid* is zero, the priority of the calling process is changed.

The *sched\_priority* member in *param* must lie between the minimum and maximum values returned by *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()*.

By default, the process priority and scheduling algorithm are inherited from or explicitly set by the parent process. Once running, the child process may change its priority by using this function.

## Returns:

- 0 Success
- 1 An error occurred (*errno* is set).

## Errors:

- EFAULT A fault occurred trying to access the buffers provided.
- EINVAL The priority isn't a valid priority.
- EPERM The calling process doesn't have sufficient privilege to set the priority.
- ESRCH The process *pid* doesn't exist.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Currently, the implementation of *sched\_setparam()* isn't 100% POSIX 1003.1-1996. The *sched\_setparam()* function sets the scheduling parameters for thread 1 in the process *pid*, or for the *calling thread* if *pid* is 0.

If you depend on this in new code, *it will not be portable*. POSIX 1003.1 says *sched\_setparam()* should return -1 and set *errno* to EPERM in a multithreaded application.



**See also:**

*errno*, *getprio()*, *sched\_getparam()*, *sched\_get\_priority\_max()*,  
*sched\_get\_priority\_min()*, *sched\_getscheduler()*, **sched\_param**,  
*sched\_setscheduler()*, *sched\_yield()*, *setprio()*

## ***sched\_setscheduler()***

© 2004, QNX Software Systems Ltd.

*Change the priority and scheduling policy of a process*

### **Synopsis:**

```
#include <sched.h>

int sched_setscheduler(
 pid_t pid,
 int policy,
 const struct sched_param *param);
```

### **Arguments:**

*pid*      The ID of the process whose priority and scheduling policy you want to set, or zero if you want to set them for the current process.

*policy*    The scheduling policy, which must be one of:

- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority ready thread runs until it blocks or is preempted by a higher priority thread.
- SCHED\_RR — similar to SCHED\_FIFO, except that threads at the same priority level timeslice (round robin) every 50 msec.
- SCHED\_OTHER — currently the same as SCHED\_RR.



Currently, you can set a thread's scheduling policy to SCHED\_SPORADIC only when you create the thread. If you use sporadic scheduling, you can't change the policy later. For more information, see *pthread\_attr\_setschedpolicy()*.

---

*param*    A pointer to a **sched\_param** structure whose *sched\_priority* member holds the priority that you want to assign to the process.

**Library:**`libc`**Description:**

The *sched\_setscheduler()* function changes the priority of process *pid* to that of the *sched\_priority* member in the `sched_param` structure passed as *param*, and the scheduling policy is set to *policy*.

If *pid* is zero, the policy and priority of the calling process are set.

The *sched\_priority* member in *param* must lie between the minimum and maximum values returned by *sched\_get\_priority\_max()* and *sched\_get\_priority\_min()*.

By default, process priority and scheduling algorithm are inherited from or explicitly set by the parent process. Once running, the child process may change its priority by using this function.

**Returns:**

The previous scheduling policy, or -1 if an error occurs (*errno* is set).

**Errors:**

|        |                                                                            |
|--------|----------------------------------------------------------------------------|
| EFAULT | A fault occurred trying to access the buffers provided.                    |
| EINVAL | The priority or scheduling policy isn't a valid value.                     |
| EPERM  | The calling process doesn't have sufficient privilege to set the priority. |
| ESRCH  | The process <i>pid</i> doesn't exist.                                      |

**Examples:**

See *sched\_getparam()*.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Currently, the implementation of *sched\_setscheduler()* isn't 100% POSIX 1003.1-1996. The *sched\_setscheduler()* function sets the scheduling policy for thread 1 in the process *pid*, or for the *calling thread* if *pid* is 0.

If you depend on this in new code, *it won't be portable*. POSIX 1003.1 says *sched\_setscheduler()* should return -1 and set *errno* to EPERM in a multithreaded application.

## See also:

*errno*, *getprio()*, *sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_get\_priority\_min()*, *sched\_getscheduler()*, ***sched\_param***, *sched\_setparam()*, *sched\_yield()*, *setprio()*

### **Synopsis:**

```
#include <sched.h>

int sched_yield(void);
```

### **Library:**

libc

### **Description:**

The *sched\_yield()* function checks to see if other threads, at the same priority as that of the calling thread, are READY to run. If so, the calling thread yields to them and places itself at the end of the READY thread queue. The *sched\_yield()* function never yields to a lower priority thread.

A higher priority thread always forces a lower priority thread to yield (that is, preempt) the instant the higher priority thread becomes ready to run, without the need for the lower priority thread to give up the processor by calling the *sched\_yield()* or *SchedYield()* functions.

The *sched\_yield()* function calls the kernel function *SchedYield()*, and may be more portable across realtime POSIX systems.



---

You should avoid designing programs that contain busy wait loops. If you can't avoid them, you can use *sched\_yield()* to reduce the system load at a given priority level. Note that a thread that calls *sched\_yield()* in a tight loop will spend a great deal of time in the kernel, which will have a small effect on interrupt latency.

---

### **Returns:**

This function always succeeds and returns zero.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <sched.h>

int main(void)
{
 int i;

 for(;;) {
 /* Process something... */
 for(i = 0 ; i < 1000 ; ++i)
 fun();

 /* Yield to anyone else at the same priority */
 sched_yield();
 }
 return EXIT_SUCCESS; /* Never reached */
}

int fun()
{
 int i;

 for(i = 0 ; i < 10 ; ++i)
 i += i;

 return(i);
}
```

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*getprio(), sched\_getparam(), sched\_get\_priority\_max(),  
sched\_get\_priority\_min(), sched\_getscheduler(), sched\_setparam(),  
sched\_setscheduler(), SchedYield(), setprio(), sleep()*

## **SchedGet(), SchedGet\_r()**

© 2004, QNX Software Systems Ltd.

*Get the scheduling policy for a thread*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SchedGet(pid_t pid,
 int tid,
 struct sched_param *param);

int SchedGet_r(pid_t pid,
 int tid,
 struct sched_param *param);
```

### **Arguments:**

*pid*        0 or a process ID; see below.

*tid*        0 or a thread ID; see below.

*param*      A pointer to a **sched\_param** structure where the function can store the scheduling parameters.

### **Library:**

libc

### **Description:**

The *SchedGet()* and *SchedGet\_r()* kernel calls return the current scheduling policy and the parameters for the thread specified by *tid* in the process specified by *pid*. If *pid* is zero, the current process is used to look up a nonzero *tid*. If *pid* and *tid* are zero, then the calling thread is used.

These functions are identical except in the way they indicate errors. See the Returns section for details.

The scheduling policy is returned on success and is one of SCHED\_FIFO, SCHED\_RR, SCHED\_SPORADIC, or SCHED\_OTHER.



**Blocking states**

These calls don't block.

**Returns:**

The only difference between these functions is the way they indicate errors:

*SchedGet()*      The current scheduling policy. If an error occurs, -1 is returned and *errno* is set.

*SchedGet\_r()*    The current scheduling policy. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

**Errors:**

EFAULT      A fault occurred when the kernel tried to access the buffers provided.

ESRCH      The process indicated by *pid* or thread indicated by *tid* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`sched_param`, *SchedInfo()*, *SchedSet()*, *SchedYield()*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SchedInfo(pid_t pid,
 int policy,
 struct _sched_info* info);

int SchedInfo_r(pid_t pid,
 int policy,
 struct _sched_info* info);
```

### **Arguments:**

- pid* A process ID, or 0 to get information about the current process.
- policy* One of the following:
- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority, ready thread runs until it blocks or is preempted by a higher-priority thread.
  - SCHED\_RR — the same as SCHED\_FIFO, except threads at the same priority level time slice (round robin) every 50 msec.
  - SCHED\_OTHER — currently the same as SCHED\_RR.
  - SCHED\_SPORADIC — sporadic scheduling. For more information, see *pthread\_attr\_setschedpolicy()*, and “Scheduling algorithms” in the chapter on the Neutrino microkernel in the *System Architecture* guide.
- info* A pointer to a `_sched_info` structure where the function can store the scheduler information.

### **Library:**

`libc`

## Description:

These kernel calls return information about the kernel's thread scheduler, including the minimum and maximum thread priority, for the process ID specified by *pid* when using the specified scheduling *policy*. If *pid* is 0, the scheduler information for the current process is returned. In either case, the `struct _sched_info` pointed to by *info* is filled in with the appropriate information.

The `SchedInfo()` and `SchedInfo_r()` functions are identical except in the way they indicate errors. See the Returns section for details.

The `struct _sched_info` structure pointed to by *info* contains at least these members:

`uint64_t interval`

The current execution time limit before the thread is suspended in favor of the scheduler.

`int priority_max`

The maximum priority for a thread using this scheduling *policy*.

`int priority_min`

The minimum priority for a thread using this scheduling *policy*.

## Returns:

The only difference between these functions is the way they indicate errors:

`SchedInfo()` If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

`SchedInfo_r()` EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

**Errors:**

- EINVAL      The *pid* or *policy* is invalid.
- ENOSYS     The *SchedInfo()* function isn't supported by this system.
- ESRCH      The process specified by *pid* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

- Cancellation point    No
- Interrupt handler     No
- Signal handler        Yes
- Thread                 Yes

**See also:**

*SchedGet()*, *SchedSet()*

## **SchedSet(), SchedSet\_r()**

© 2004, QNX Software Systems Ltd.

*Set the scheduling policy for a thread*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SchedSet(
 pid_t pid,
 int tid,
 int policy,
 const struct sched_param *param);

int SchedSet_r(
 pid_t pid,
 int tid,
 int policy,
 const struct sched_param *param);
```

### **Arguments:**

- pid*      0 or a process ID; see below.
- tid*      0 or a thread ID; see below.
- policy*    The scheduling policy; one of:
- SCHED\_FIFO — a fixed-priority scheduler in which the highest priority, ready thread runs until it blocks or is preempted by a higher priority thread.
  - SCHED\_RR — the same as SCHED\_FIFO, except threads at the same priority level time slice (round robin) every 50 msec.
  - SCHED\_OTHER — currently the same as SCHED\_RR.
  - SCHED\_NOCHANGE — this isn't actually a policy, but a special value that tells the kernel to update the parameters specified in *param*, without changing the policy.



---

Currently, you can set a thread's scheduling policy to SCHED\_SPORADIC only when you create the thread. If you use sporadic scheduling, you can't change the policy later. For more information, see *pthread\_attr\_setschedpolicy()*.

---

*param* A pointer to a **sched\_param** structure where the function can store the scheduling parameters.

## Library:

**libc**

## Description:

The *SchedSet()* and *SchedSet\_r()* kernel calls set both the scheduling policy and the associated parameters for the thread specified by *tid* in the process specified by *pid*. If *pid* is zero the current process is used to look up a nonzero *tid*. If *tid* is zero, then the calling thread is used and *pid* is ignored.

These functions are identical except in the way they indicate errors. See the Returns section for details.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

*SchedSet()* If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*SchedSet\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

|        |                                                                                       |
|--------|---------------------------------------------------------------------------------------|
| EFAULT | A fault occurred when the kernel tried to access the buffers you provided.            |
| EINVAL | The given scheduling <i>policy</i> is invalid.                                        |
| EPERM  | The process doesn't have permission to change the scheduling of the indicated thread. |
| ESRCH  | The process indicated by <i>pid</i> or thread indicated by <i>tid</i> doesn't exist.  |

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*sched\_get\_priority\_max()*, *sched\_get\_priority\_min()*, **sched\_param**, *SchedGet()*, *SchedInfo()*, *SchedYield()*



**Synopsis:**

```
#include <sys/neutrino.h>

int SchedYield(void);

int SchedYield_r(void);
```

**Library:**

```
libc
```

**Description:**

These kernel calls check to see if other threads at the same priority as that of the calling thread are ready to run. If so, the calling thread yields to them and places itself at the end of the ready thread queue for that priority. *SchedYield()* never yields to a lower priority thread. Higher priority threads always force a yield the instant they become ready to run. This call has no effect with respect to threads running at priorities other than the calling thread's.

The *SchedYield()* and *SchedYield\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

Avoid designing programs that contain busy-wait loops using *SchedYield()* to timeslice. If this is unavoidable, you can use *SchedYield\_r()* to reduce the system load at a given priority level. Note that a program that calls *SchedYield()* in a tight loop will spend a great deal of time in the kernel, which will have a small effect on scheduling interrupt latency.

**Blocking states**

These calls don't block. However, if other threads are ready at the same priority, the calling thread is placed at the end of the ready queue for this priority.

## Returns:

The only difference between these functions is the way they indicate errors:

*SchedYield()* If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*SchedYield\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*SchedGet()*, *SchedSet()*

*Add or remove one or more given addresses from an association***Synopsis:**

```
#include <netinet/sctp.h>

int sctp_bindx(int sd,
 struct sockaddr *addrs,
 int addrcnt,
 int flags);
```

**Arguments:**

- sd* Socket descriptor. Depending on the type of *sd*, the type of address is determined. If *sd* is an **IPv4** socket, the address passed is an **IPv4** address. If *sd* is an **IPv6** socket, the address passed is either an **IPv4** or an **IPv6** address. A single address is specified as **INADDR\_ANY** or **IN6ADDR\_ANY**.
- addrs* Pointer to an array of one or more socket addresses. Each address is contained in its appropriate structure (i.e. struct **sockaddr\_in** or struct **sockaddr\_in6**). The family of the address type must be used to distinguish the address length. This representation is termed a “packed array” of addresses.
- addrcnt* Number of addresses in the array.
- flags* Either **SCTP\_BINDX\_ADD\_ADDR** or **SCTP\_BINDX\_REM\_ADDR**. It is formed from the bitwise OR of zero or more of these flags.

**Library:**

**libsctp**

**Description:**

The *sctp\_bindx()* add or remove one or more given addresses from an association.

An application uses *sctp\_bindx(SCTP\_BINDX\_ADD\_ADDR)* to associate additional addresses with an endpoint after calling *bind()*. Otherwise, it may also call *sctp\_bindx(SCTP\_BINDX\_REM\_ADDR)* to remove some addresses a listening socket is associated with — so that no new association will be associated with those addresses. If the endpoint supports dynamic address, a *SCTP\_BINDX\_REM\_ADDR* or *SCTP\_BINDX\_ADD\_ADDR* may cause an endpoint to send the appropriate message to change the peer's address lists.

**Returns:**

- 0 Success.
- 1 Failure; *errno* is set.

**Errors:**

- EFAULT Passed-in flag was neither *SCTP\_BINDX\_ADD\_ADDR* nor *SCTP\_BINDX\_REM\_ADDR*.
- EINVAL Passed-in address has a wrong family.

**Classification:**

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

SCTP, *sctp\_connectx()*, *sctp\_freeladdrs()*, *sctp\_freepaddrs()*,  
*sctp\_getladdrs()*, *sctp\_getpaddrs()*, *sctp\_peeloff()*, *sctp\_rcvmsg()*,  
*sctp\_sendmsg()*

## ***sctp\_connectx()***

© 2004, QNX Software Systems Ltd.

*Help associate an endpoint that is multi-homed*

### **Synopsis:**

```
#include <netinet/sctp.h>

int sctp_connectx(int s,
 struct sockaddr *addrs,
 int addrcnt);
```

### **Arguments:**

*s*            Socket descriptor.

*addrs*        Array of addresses.

*addrcnt*      Number of addresses in the array.

### **Library:**

`libsctp`

### **Description:**

The *sctp\_connectx()* function connects a host to a multi-homed endpoint by specifying a list of peer addresses.

### **Returns:**

0            Success.

-1          Failure; *errno* is set.

### **Errors:**

EINVAL      No valid address is passed in *addrs*.

### **Classification:**

Socket API extension for stream control transmission protocol in accordance with *draft-ietf-tsvwg-sctpsocket-07.txt*.

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

SCTP, *sctp\_bindx()*, *sctp\_freeladdrs()*, *sctp\_freepaddrs()*,  
*sctp\_getladdrs()*, *sctp\_getpaddrs()*, *sctp\_peeloff()*, *sctp\_rcvmsg()*,  
*sctp\_sendmsg()*

## ***sctp\_freeladdr()***

© 2004, QNX Software Systems Ltd.

*Free all resources allocated by sctp\_getladdr()*

### **Synopsis:**

```
#include <netinet/sctp.h>

void sctp_freeladdr(struct sockaddr *addr);
```

### **Arguments:**

*addr*     Array of peer addresses returned by *sctp\_getladdr()*.

### **Library:**

`libsctp`

### **Description:**

The *sctp\_freeladdr()* free all resources allocated by *sctp\_getladdr()*.

### **Classification:**

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freepaddr()*, *sctp\_getladdr()*, *sctp\_getpaddr()*, *sctp\_peeloff()*, *sctp\_rcvmsg()*, *sctp\_sendmsg()*



### **Synopsis:**

```
#include <netinet/sctp.h>

void sctp_freepaddr(struct sockaddr *addr);
```

### **Arguments:**

*addr*     Array of peer addresses returned by *sctp\_getpaddr()*.

### **Library:**

libsctp

### **Description:**

The *sctp\_freepaddr()* free all resources allocated by *sctp\_getpaddr()*.

### **Classification:**

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdr()*, *sctp\_getladdr()*, *sctp\_getpaddr()*, *sctp\_peeloff()*, *sctp\_rcvmsg()*, *sctp\_sendmsg()*

## ***sctp\_getladdr***(***s***)

© 2004, QNX Software Systems Ltd.

*Return all locally bound addresses on a socket*

### **Synopsis:**

```
#include <netinet/sctp.h>

int sctp_getladdr(int sd,
 sctp_assoc_t id,
 struct sockaddr **addr);
```

### **Arguments:**

*sd*        Socket descriptor.

*id*        Specifies the association for one-to-many style sockets. It is ignored for one-to-one style sockets.

*addr*      A pointer to an array of local addresses, returned by the stack.

### **Library:**

`libsctp`

### **Description:**

The *sctp\_getladdr*(*s*) return all locally bound addresses on a socket.

On return, *addr* points to a dynamically allocated packed array of *sockaddr* structures of the appropriate type for each address. The caller should use *sctp\_freepaddr*(*s*) to free the memory. Note that the in/out parameter *addr* must not be NULL.

### **Returns:**

On success, *sctp\_getladdr*(*s*) returns the number of local addresses bound to the socket. If the socket is unbound, *sctp\_getladdr*(*s*) returns 0, and the value of *addr* is undefined.

If an error occurs, *sctp\_getladdr*(*s*) returns -1, and the value of *addr* is undefined.

## Errors:

|          |                                    |
|----------|------------------------------------|
| EINVAL   | Address is wrong.                  |
| ENOTCONN | Socket is not bound.               |
| ENOMEM   | Can't allocate memory for address. |

## Classification:

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdr()*, *sctp\_freepaddr()*, *sctp\_getpaddr()*, *sctp\_peeloff()*, *sctp\_recvmg()*, *sctp\_sendmsg()*

## ***sctp\_getpaddr***(***s***)

© 2004, QNX Software Systems Ltd.

*Return all peer addresses in an association*

### **Synopsis:**

```
#include <netinet/sctp.h>

int sctp_getpaddr(int sd,
 sctp_assoc_t id,
 struct sockaddr **addr);
```

### **Arguments:**

*sd*        Socket descriptor.

*id*        Specifies the association for one-to-many style sockets. It is ignored for one-to-one style sockets

*addr*      A pointer to an array of peer addresses, returned by the stack.

### **Library:**

`libsctp`

### **Description:**

The *sctp\_getpaddr*(*s*) return all peer addresses in an association.

On return, *addr* points to a dynamically packed array of *sockaddr* structures of the appropriate type for each address. The caller should use *sctp\_freepaddr*(*s*) to free the memory. Note that in and out parameters *addr* must not be NULL.

### **Returns:**

On success, *sctp\_getpaddr*(*s*) returns the number of peer addresses in the association. If there is no association, this function returns 0 and the value of *\*addr* is undefined.

If an error occurs, *sctp\_getpaddr*(*s*) returns -1, and the value of *\*addr* is undefined.

## Errors:

- |        |                                    |
|--------|------------------------------------|
| EINVAL | Passed-in address is wrong.        |
| ENOMEM | Can't allocate memory for address. |

## Classification:

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdr()*, *sctp\_freepaddr()*, *sctp\_getladdr()*, *sctp\_peeloff()*, *sctp\_rcvmsg()*, *sctp\_sendmsg()*

## ***sctp\_peeloff()***

© 2004, QNX Software Systems Ltd.

*Branch off an association into a separate socket*

### **Synopsis:**

```
#include <netinet/sctp.h>

int sctp_peeloff(int sd,
 sctp_assoc_t assoc_id);
```

### **Library:**

libsctp

### **Description:**

You call this function to branch off an association into a separate socket. The new socket is a one-to-one style socket. You should confine your operation to one that is allowed for a one-to-one style socket.

Using *sctp\_peeloff()*, you create a new socket descriptor as follows:

```
new_sd = sctp_peeloff(int sd, sctp_assoc_t assoc_id);
```

### **Returns:**

On success, it returns a new socket, which has the single association in it. On failure, it returns -1 and *errno* is set.

### **Errors:**

|          |                                   |
|----------|-----------------------------------|
| EBADF    | The <i>sd</i> is wrong.           |
| ENOTSOCK | Can't branch off the association. |

### **Classification:**

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdrs()*,  
*sctp\_freepaddrs()*, *sctp\_getladdrs()*, *sctp\_getpaddrs()*, *sctp\_recvmmsg()*,  
*sctp\_sendmmsg()*

# SCTP

© 2004, QNX Software Systems Ltd.

*Stream Control Transmission Protocol*

## Synopsis:

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(PF_INET,
 SOCK_DGRAM,
 IPPROTO_SCTP);
```

Or,

```
int socket(PF_INET,
 SOCK_STREAM,
 IPPROTO_SCTP);
```

## Description:

The SCTP protocol provides reliable end-to-end message transport service. It has the following features:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to path MTU size
- Sequenced delivery of user messages within multiple streams
- Multi-homing
- Protection against DOS attacks.

## Returns:

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

## Errors:

|        |                                                                                |
|--------|--------------------------------------------------------------------------------|
| EACCES | Permission to create a socket of the specified type and/or protocol is denied. |
| EMFILE | The per-process descriptor table is full.                                      |



|                 |                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------|
| ENFILE          | The system file table is full.                                                                         |
| ENOBUFS         | Insufficient buffer space available. The socket can't be created until sufficient resources are freed. |
| ENOMEM          | Not enough memory.                                                                                     |
| EPROTONOSUPPORT | The protocol type or the specified protocol isn't supported within this domain.                        |

### See also:

#### IP

*sctp\_bindx(), sctp\_connectx(), sctp\_freeladdrs(), sctp\_freepaddrs(), sctp\_getladdrs(), sctp\_getpaddrs(), sctp\_peeloff(), sctp\_recvmmsg(), sctp\_sendmmsg()*

*RFC 2960, RFC 3257*

#### Drafts:

- Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.
- Stream control transmission protocol dynamic address reconfiguration.

## **sctp\_recvmsg()**

© 2004, QNX Software Systems Ltd.

*Receive message using advanced SCTP features*

### **Synopsis:**

```
#include <netinet/sctp.h>

ssize_t sctp_recvmsg(int s,
 void *msg,
 size_t len,
 struct sockaddr *from,
 socklen_t *fromlen,
 struct sctp_sndrcvinfo *sinfo,
 int *msg_flags);
```

### **Arguments:**

|                  |                                                                                                                                                              |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>         | Socket descriptor.                                                                                                                                           |
| <i>msg</i>       | Message buffer to be filled.                                                                                                                                 |
| <i>len</i>       | Length of the message buffer.                                                                                                                                |
| <i>from</i>      | A pointer to a <b>sockaddr</b> object where the function can store the source address of the message.                                                        |
| <i>fromlen</i>   | A pointer to a <b>socklen_t</b> object that specifies the size of the <i>from</i> buffer. The function stores the actual size of the address in this object. |
| <i>sinfo</i>     | A pointer to a <b>sctp_sndrcvinfo</b> structure to be filled upon receipt of the message.                                                                    |
| <i>msg_flags</i> | A pointer to an integer to be filled with any message flags (e.g. MSG_NOTIFICATION).                                                                         |

### **Library:**

**libsctp**

**Description:**

The *sctp\_rcvmsg()* function receive a message from socket *s*, whether or not a socket is connected. The difference between this function and the generic function *rcvmsg()* is that you could pass in a pointer to **sctp\_sndrcvinfo** structure, and the structure is filled upon receipt of the message. The structure has detailed information about the message you just received. Note that you must enable **sctp\_data\_io\_events** with the Sctp\_EVENTS socket option first, to be able to have the **sctp\_sndrcvinfo** structure be filled in.

**Returns:**

Number of bytes received, or -1 if an error occurs (*errno* is set).

**Errors:**

ENOMEM      Not enough stack memory.

**Classification:**

Socket API extension for stream control transmission protocol in accordance with *draft-ietf-tsvwg-sctpsocket-07.txt*.

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdrs()*, *sctp\_freepaddrs()*, *sctp\_getladdrs()*, *sctp\_getpaddrs()*, *sctp\_peeloff()*, *sctp\_sendmsg()*

## **sctp\_sendmsg()**

© 2004, QNX Software Systems Ltd.

*Send message using advanced SCTP features*

### **Synopsis:**

```
#include <netinet/sctp.h>

ssize_t sctp_sendmsg(int s,
 const void *msg,
 size_t len,
 struct sockaddr *to,
 socklen_t tolen,
 uint32_t ppid,
 uint32_t flags,
 uint16_t stream_no,
 uint32_t timetolive,
 uint32_t context);
```

### **Arguments:**

|              |                                                                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>     | Socket descriptor.                                                                                                                                                                                               |
| <i>msg</i>   | Message to be sent.                                                                                                                                                                                              |
| <i>len</i>   | Length of the message.                                                                                                                                                                                           |
| <i>to</i>    | Destination address of the message.                                                                                                                                                                              |
| <i>tolen</i> | Length of the destination address.                                                                                                                                                                               |
| <i>ppid</i>  | An opaque unsigned value that is passed to the remote end in each user message. The byte order issues are not accounted for and this information is passed opaquely by the SCTP stack from one end to the other. |
| <i>flags</i> | Flags composed of bitwise OR of these values:<br><br>MSG_UNORDERED<br><br>This flag requests the un-ordered delivery of the message. If the flag is clear, the datagram is considered an ordered send.           |

**MSG\_ADDR\_OVER**

This flag, in one-to-many style, requests the SCTP stack to override the primary destination address.

**MSG\_ABORT** This flag causes the specified association to abort — by sending an ABORT message to the peer (one-to-many style only).

**MSG\_EOF** This flag invokes the SCTP graceful shutdown procedures on the specified association. Graceful shutdown assures that all data enqueued by both endpoints is successfully transmitted before closing the association (one-to-many style only).

*stream\_no* Message stream number — for the application to send a message. If a sender specifies an invalid stream number, an error indication is returned and the call fails.

*timetolive* Message time to live in milliseconds. The sending side expires the message within the specified time period if the message has not been sent to the peer within this time period. This value overrides any default value set using socket option. If you use a value of 0, it indicates that no timeout should occur on this message.

*context* An opaque 32-bit context datum. This value is passed back to the upper layer if an error occurs while sending a message, and is retrieved with each undelivered message.

**Library:**

**libsctp**

**Description:**

The *sctp\_sendmsg()* function allows you to send extra information to a remote application. Using advanced SCTP features, you can send a message through a specified stream, pass extra opaque information to a remote application, or define a timeout for the particular message.

**Returns:**

The number of bytes sent, or -1 if an error occurs (*errno* is set).

**Errors:**

|              |                                                                                                            |
|--------------|------------------------------------------------------------------------------------------------------------|
| EBADF        | An invalid descriptor was specified.                                                                       |
| EDESTADDRREQ | A destination address is required.                                                                         |
| EFAULT       | An invalid user space address was specified for a parameter.                                               |
| EMSGSIZE     | The socket requires that the message be sent atomically, but the size of the message made this impossible. |
| ENOBUFS      | The system couldn't allocate an internal buffer. The operation may succeed when buffers become available.  |
| ENOTSOCK     | The argument <i>s</i> isn't a socket.                                                                      |
| EWOULDBLOCK  | The socket is marked nonblocking and the requested operation would block.                                  |

**Classification:**

Socket API extension for stream control transmission protocol in accord with *draft-ietf-tsvwg-sctpsocket-07.txt*.

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

SCTP, *sctp\_bindx()*, *sctp\_connectx()*, *sctp\_freeladdrs()*,  
*sctp\_freepaddrs()*, *sctp\_getladdrs()*, *sctp\_getpaddrs()*, *sctp\_peeloff()*,  
*sctp\_rcvmsg()*,

# **searchenv()**

© 2004, QNX Software Systems Ltd.

*Search the directories listed in an environment variable*

## **Synopsis:**

```
#include <stdlib.h>


void searchenv(const char* name,
 const char* env_var,
 char* buffer);
```

## **Arguments:**

*name*            The name of the file that you want to search for.

*env\_var*         The name of an environment variable whose value is a list of directories that you want to search. Common values for *env\_var* are "**PATH**", "**LIB**" and "**INCLUDE**".

---

 The *searchenv()* function doesn't search the current directory unless it's specified in the environment variable.

---

*buffer*          A buffer where the function can store the full path of the file found. This buffer should be `PATH_MAX` bytes long. If the specified file can't be found, the function stores an empty string in the buffer.


## **Library:**

`libc`

## **Description:**

The *searchenv()* function searches for the file specified by *name* in the list of directories assigned to the environment variable specified by *env\_var*.

---

 Use *pathfind()* or *pathfind\_r()* instead of this function.

---



**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <limits.h>

void display_help(FILE *fp)
{
 printf("display_help T.B.I.\n");
}

int main(void)
{
 FILE *help_file;
 char full_path[PATH_MAX];

 searchenv("lib_ref.html", "PATH", full_path);
 if(full_path[0] == '\0') {
 printf("Unable to find help file\n");
 } else {
 help_file = fopen(full_path, "r");
 display_help(help_file);
 fclose(help_file);
 }

 return EXIT_SUCCESS;
}
```

**Classification:**

QNX 4

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **Caveats:**

The *searchenv()* function manipulates the environment pointed to by the global *environ* variable.

### **See also:**

*getenv()*, *pathfind()*, *pathfind\_r()*, *setenv()*

**Synopsis:**

```
#include <stdlib.h>

unsigned short int *seed48(
 unsigned short int seed16v[3]);
```

**Arguments:**

*seed16v*     An array that comprises the 48 bits of the seed.

**Library:**

`libc`

**Description:**

The *seed48()* initializes the internal buffer *r(n)* of *drand48()*, *lrand48()*, and *mrnd48()*. All 48 bits of the seed can be specified in an array of 3 short integers, where the entry with index 0 specifies the lowest bits. The constant multiplicand and addend of the algorithm are reset to the defaults: the multiplicand  $a = 0xFDEECE66D = 25214903917$  and the addend  $c = 0xB = 11$ .

**Returns:**

A pointer to an array of 3 shorts which contains the old seed. This array is statically allocated, thus its contents are lost after each new call to *seed48()*.

**Classification:**

Standard Unix

**Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*drand48(), erand48(), jrand48(), lcong48(), lrand48(), mrand48(),  
nrand48(), srand48()*

## Synopsis:

```
#include <dirent.h>

void seekdir(DIR * dirp,
 long int pos);
```

## Arguments:

*dirp*      A pointer to the directory stream, for which you want to set the current location.

*pos*        The new position for the directory stream. You should have obtained this value from an earlier call to *telldir()*.

## Library:

libc

## Description:

The *seekdir()* function sets the position of the next *readdir()* operation on the directory stream specified by *dirp* to the position specified by *pos*.

The new position reverts to the one associated with the directory stream when the *telldir()* operation was performed.

Values returned by *telldir()* are good only for the lifetime of the **DIR** pointer, *dirp*, from which they're derived. If the directory is closed and then reopened, the *telldir()* value may be invalidated due to undetected directory compaction. It's safe to use a previous *telldir()* value immediately after a call to *opendir()* and before any calls to *readdir()*.

## Classification:

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*closedir(), errno, lstat(), opendir(), readdir(), readdir\_r(), rewinddir(),  
telldir(), stat()*

**Synopsis:**

```
#include <sys/select.h>

int select(int width,
 fd_set * readfds,
 fd_set * writefds,
 fd_set * exceptfds,
 struct timeval * timeout);

FD_SET(int fd, fd_set * fdset);
FD_CLR(int fd, fd_set * fdset);
FD_ISSET(int fd, fd_set * fdset);
FD_ZERO(fd_set * fdset);
```

**Arguments:**

- width*            The number of descriptors to check in the given sets. Only the descriptors from 0 through (*width*-1) in the descriptor sets are examined. Therefore, the value of *width* must be at least as large as:
- (highest valued file descriptor in the sets) + 1*
- readfds*        NULL, or a pointer to a **fd\_set** object that specifies the descriptors to check for files that are ready for reading. The function replaces the set with the file descriptors that are actually ready for reading.
- writefds*       NULL, or a pointer to a **fd\_set** object that specifies the descriptors to check for files that are ready for writing. The function replaces the set with the file descriptors that are actually ready for writing.
- exceptfds*      NULL, or a pointer to a **fd\_set** object that specifies the descriptors to check for files that have an exceptional condition pending. The function replaces the set with the file descriptors that actually have an exceptional condition pending.

*timeout* NULL, or a pointer to a `timeval` that specifies how long to wait for the selection to complete.

## Library:

`libc`

## Description:

The *select()* function examines the file descriptor sets whose addresses are passed in *readfds*, *writefds*, and *exceptfds* to see if some of their descriptors are ready for reading, ready for writing, or have an exceptional condition pending. Any of *readfds*, *writefds*, and *exceptfds* may be NULL pointers if no descriptors are of interest.



---

In earlier versions of QNX Neutrino, *select()* and the associated macros were defined in `sys/time.h`. They're now defined in `sys/select.h`, which `sys/time.h` includes.

---

The *select()* function replaces the given descriptor sets with subsets consisting of those descriptors that are ready for the requested operation, and returns the total number of ready descriptors in all the sets.

If *timeout* isn't NULL, it specifies a maximum interval to wait for the selection to complete. If *timeout* is NULL, *select()* blocks until one of the selected conditions occurs. To effect a poll, the *timeout* argument should be a non-NULL pointer, pointing to a zero-valued `timeval` structure.

If the current operating system configuration supports a larger number of open files than is specified in `FD_SETSIZE`, you can increase the number of open file descriptors used with *select()* by changing the definition of `FD_SETSIZE` before including `<sys/select.h>` or `<sys/time.h>`.

If you use *select()* with a timeout, you should reset the timeout value after calling *select()*.





---

If you're using *select()* in conjunction with the socket API package, note that selecting for reading on a socket descriptor on which a *listen()* has been performed indicates that a subsequent *accept()* on that descriptor won't block.

---

### Manipulating file-descriptor sets

At least the following macros are defined in `<sys/select.h>` for manipulating file-descriptor sets:

*FD\_ZERO*( *&fdset* )

Initialize a descriptor set *fdset* to the null set.

*FD\_SET*( *fd*, *&fdset* )

Add the file descriptor *fd* to the set *fdset*.

*FD\_CLR*( *fd*, *&fdset* )

Remove *fd* from *fdset*.

*FD\_ISSET*( *fd*, *&fdset* )

Is nonzero if *fd* is a member of *fdset*; otherwise, zero.

The behavior of these macros is undefined if a descriptor value is less than zero, or greater than or equal to `FD_SETSIZE`.

### Returns:

The number of ready descriptors in the descriptor sets, 0 if the *timeout* expired, or -1 if an error occurs (*errno* is set).

### Errors:

|        |                                                                                                               |
|--------|---------------------------------------------------------------------------------------------------------------|
| EBADF  | One of the descriptor sets specified an invalid descriptor.                                                   |
| EFAULT | One of the pointers given in the call referred to a nonexistent portion of the address space for the process. |

- EINTR**      A signal was delivered before any of the selected events occurred, or before the time limit expired.
- EINVAL**     A component of the pointed-to time limit is outside the acceptable range: *t\_sec* must be between 0 and  $10^8$ , inclusive; *t\_usec* must be greater than or equal to 0, and less than  $10^6$ .

**Examples:**

```
/*
 * This example opens a console and a serial port for
 * read mode, and calls select() with a 5 second timeout.
 * It waits for data to be available on either descriptor.
 */

#include <unistd.h>
#include <stdlib.h>
#include <fcntl.h>
#include <sys/select.h>

int main(void)
{
 int console, serial;
 struct timeval tv;
 fd_set rfd;
 int n;

 if((console = open("/dev/con1", O_RDONLY)) == -1
 || (serial = open("/dev/ser1", O_RDONLY)) == -1)
 {
 perror("open");
 return EXIT_FAILURE;
 }

 /*
 * Clear the set of read file descriptors, and
 * add the two we just got from the open calls.
 */
 FD_ZERO(&rfd);
 FD_SET(console, &rfd);
 FD_SET(serial, &rfd);

 /*
 * Set a 5 second timeout.
 */
 tv.tv_sec = 5;
 tv.tv_usec = 0;
```

```
switch (n = select(1 + max(console, serial),
 &rfd, 0, 0, &tv)) {
case -1:
 perror("select");
 return EXIT_FAILURE;
case 0:
 puts("select timed out");
 break;
default:
 printf("%d descriptors ready ...\n", n);
 if(FD_ISSET(console, &rfd))
 puts(" -- console descriptor has data pending");
 if(FD_ISSET(serial, &rfd))
 puts(" -- serial descriptor has data pending");
}
return EXIT_SUCCESS;
}
```

## Classification:

Standard Unix

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## Caveats:

The *select()* function only works with raw file descriptors; it doesn't work with file descriptors in edited mode. See the ICANON flag in the description of the *tcgetattr()* function.

**See also:**

*errno, fcntl(), read(), sysconf(), tcsetattr(), write()*

**Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int select_attach
(void *dpp,
 select_attr_t *attr,
 int fd,
 unsigned flags,
 int (*func) (select_context_t *ctp,
 int fd,
 unsigned flags,
 void *handle),
 void *handle);
```

**Arguments:**

*dpp* The dispatch handle, as returned by *dispatch\_create()*, that you want to attach to a file descriptor.

*attr* A pointer to a **select\_attr\_t** structure. This structure is defined as:

```
typedef struct _select_attr {
 unsigned flags;
} select_attr_t;
```

Currently, no attribute flags are defined.

*fd* The file descriptor that you want to attach to the dispatch handle.

*flags* Flags that specify the events that you're interested in. For more information, see "Flags," below.

*func* The function that you want to call when the file descriptor unblocks. For more information, see "Function," below.

*handle* A pointer to arbitrary data that you want to pass to *func*.

## Library:

`libc`

## Description:

The function *select\_attach()* attaches the file descriptor *fd* to the dispatch handle *dpp* and selects *flags* events. When *fd* “unblocks”, *func* is called with *handle*.

## Flags

The available flags are defined in `<sys/dispatch.h>`. The following flags use *ionotify()* mechanisms (see *ionotify()* for further details):

### SELECT\_FLAG\_EXCEPT

Out-of-band data is available. The definition of out-of-band data depends on the resource manager.

### SELECT\_FLAG\_READ

There's input data available. The amount of data available defaults to 1. For a character device such as a serial port, this is a character. For a POSIX message queue, it's a message. Each resource manager selects an appropriate object.

### SELECT\_FLAG\_WRITE

There's room in the output buffer for more data. The amount of room available needed to satisfy this condition depends on the resource manager. Some resource managers may default to an empty output buffer, while others may choose some percentage of the empty buffer.

These flags are specific to dispatch:

### SELECT\_FLAG\_REARM

Rearm the *fd* after an event is dispatched.

**SELECT\_FLAG\_SRVEXCEPT**

Register a function that's called whenever a server, to which this client is connected, dies. (This flag uses the *ChannelCreate()* function's `_NTO_CHF_COID_THREADDEATH` flag. In this case, *fd* is ignored.)

**Function**

The argument *func* is the user-supplied function that's called when one of the registered events occurs on *fd*. This function should return 0 (zero); other values are reserved. The function is passed the following arguments:

*ctp* Context pointer.

*fd* The *fd* on which the event occurred.

*flags* The type of event that occurred. The possible *flags* are:

- SELECT\_FLAG\_EXCEPT
- SELECT\_FLAG\_READ
- SELECT\_FLAG\_WRITE

For descriptions of the flags passed to *func*, see "Flags," above.

*handle* The *handle* passed to *select\_attach()*.

**Returns:**

Zero on success, or -1 if an error occurred (*errno* is set).

**Errors:**

EINVAL Invalid argument.

ENOMEM Insufficient memory was available.

## Examples:

For an example with *select\_attach()*, see *dispatch\_create()*. For other examples using the dispatch interface, see *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*select\_detach()*, *select\_query()*



### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int select_detach(void *dpp,
 int fd);
```

### **Arguments:**

*dpp*     The dispatch handle, as returned by *dispatch\_create()*, that you want to detach from the file descriptor.

*fd*      The file descriptor that you want to detach.

### **Library:**

libc

### **Description:**

The function *select\_detach()* detaches the file descriptor *fd* that was registered with dispatch *dpp*, using the *select\_attach()* call.

### **Returns:**

0        Success.

-1      The file descriptor *fd* wasn't registered with the dispatch *dpp*.

### **Examples:**

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>
#include <fcntl.h>

int my_func(...) {
 :
}

int main(int argc, char **argv) {
```

```
dispatch_t *dpp;
int fd;
select_attr_t attr;

if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
}

if(argc <= 2 || (fd = open(argv[1],
 O_RDWR | O_NONBLOCK)) == -1) {
 exit(0);
}

select_attach(dpp, &attr, fd,
 SELECT_FLAG_READ | SELECT_FLAG_REARM, my_func, NULL);

:

if ((select_detach(dpp, fd)) == -1) {
 fprintf(stderr, "Failed to detach \
 the file descriptor.\n");
 return 1;
}
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*select\_attach(), select\_query()*

## ***select\_query()***

© 2004, QNX Software Systems Ltd.

*Decode the last select event*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int select_query
(select_context_t *ctp,
 int *fd,
 unsigned *flags,
 int (**func) (select_context_t *ctp,
 int fd,
 unsigned flags,
 void *handle),
 void **handle);
```

### **Arguments:**

|               |                                                                                                                                                                     |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ctp</i>    | A pointer to a <code>select_context_t</code> structure that defines the context of the event that you want to get information about.                                |
| <i>fd</i>     | A pointer to a location where the function can store the file descriptor that's associated with the event.                                                          |
| <i>flags</i>  | A pointer to a location where the function can store the flags associated with the event; see "Flags" in the documentation for <code>select_attach()</code> .       |
| <i>func</i>   | A pointer to a location where the function can store the function associated with the event; see "Function" in the documentation for <code>select_attach()</code> . |
| <i>handle</i> | A pointer to a location where the function can store the address of any data that you arranged to pass to <code>func</code> .                                       |

### **Library:**

`libc`

## Description:

The function *select\_query()* stores the values of the last select event, for context *ctp*, in *fd*, *flags*, *func*, and *handle*.

## Returns:

If an error occurs, the function returns -1. An error occurs if the received event doesn't belong to one of the file descriptors attached with *select\_attach()*.

## Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>
#include <fcntl.h>

int my_func(select_context_t *ctp,
 int fd,
 unsigned flags,
 void *handle) {

 :

}

int main(int argc, char **argv) {
 dispatch_t *dpp;
 dispatch_context_t *ctp;
 int fd;
 unsigned flag;
 void *handle;
 select_attr_t *attr;
 int (*func)(select_context_t *,
 int, unsigned, void *);

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch handle.\n", argv[0]);
 return EXIT_FAILURE;
 }

 if(argc ≤ 2 || (fd = open(argv[1],
 O_RDWR | O_NONBLOCK)) == -1) {
 exit(0);
 }
}
```

```
select_attach(dpp, attr, fd,
 SELECT_FLAG_READ | SELECT_FLAG_REARM, &my_func, NULL);

ctp = dispatch_context_alloc(dpp);

:

if(select_query((select_context_t *)ctp, &fd, &flag,
 &func, &handle) == -1) {
 fprintf(stderr, "Failed to decode last select event.\n");
 return 1;
}
}
```

For more examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*select\_attach()*, *select\_detach()*

## Synopsis:

```
#include <semaphore.h>

int sem_close(sem_t * sem);
```

## Arguments:

*sem*     A pointer to a semaphore, as returned by *sem\_open()*.

## Library:

`libc`

## Description:

The *sem\_close()* function closes the named semaphore *sem* opened by *sem\_open()*, releasing any system resources associated with the *sem*.



---

Don't mix named semaphore operations (*sem\_open()* and *sem\_close()*) with unnamed semaphore operations (*sem\_init()* and *sem\_destroy()*) on the same semaphore.

---

## Returns:

0     Success.

-1    An error occurred (*errno* is set).

## Errors:

EINVAL     Invalid semaphore descriptor *sem*.

## Classification:

POSIX 1003.1

## **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **Caveats:**

The `mqueue` manager must be running for applications to use named semaphores.

## **See also:**

*sem\_init()*, *sem\_open()*, *sem\_unlink()*

`mqueue` in the *Utilities Reference*



### **Synopsis:**

```
#include <semaphore.h>

int sem_destroy(sem_t * sem);
```

### **Arguments:**

*sem* A pointer to the **sem\_t** object for the semaphore that you want to destroy.

### **Library:**

libc

### **Description:**

The *sem\_destroy()* function destroys the unnamed semaphore referred to by the *sem* argument. The semaphore must have been previously initialized by the *sem\_init()* function.

The effect of using a semaphore after it has been destroyed is undefined. If you destroy a semaphore that other processes are currently blocked on, they're unblocked, with an error (EINVAL).



---

Don't mix named semaphore operations (*sem\_open()* and *sem\_close()*) with unnamed semaphore operations (*sem\_init()* and *sem\_destroy()*) on the same semaphore.

---

### **Returns:**

0 Success.

-1 An error occurred (*errno* is set).

**Errors:**

EINVAL     Invalid semaphore descriptor *sem*.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sem\_init()*, *sem\_post()*, *sem\_trywait()*, *sem\_wait()*

### **Synopsis:**

```
#include <semaphore.h>

int sem_getvalue(sem_t * sem,
 int * value);
```

### **Arguments:**

*sem*        A pointer to the **sem\_t** object for the semaphore whose value you want to get.

*value*     A pointer to a location where the function can store the semaphore's value. A positive value (i.e. greater than zero) indicates an unlocked semaphore, and a *value* of 0 (zero) indicates a locked semaphore.

### **Library:**

**libc**

### **Description:**

The *sem\_getvalue()* function takes a snapshot of the value of the semaphore, *sem*, and stores it in *value*. This value can change at any time, and is guaranteed valid only at some point in the *sem\_getvalue()* call.

This function is provided for debugging semaphore code.

### **Returns:**

0        Success.

-1      An error occurred (*errno* is set).

## **Errors:**

EINVAL     Invalid semaphore descriptor *sem*.

## **Classification:**

POSIX 1003.1

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*sem\_destroy()*, *sem\_init()*, *sem\_trywait()*, *sem\_wait()*

## Synopsis:

```
#include <semaphore.h>

int sem_init(sem_t * sem,
 int pshared,
 unsigned value);
```

## Arguments:

|                |                                                                                                                                                                                                                       |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>sem</i>     | A pointer to the <b>sem_t</b> object for the semaphore that you want to initialize.                                                                                                                                   |
| <i>pshared</i> | Nonzero if you want the semaphore to be shared between processes via shared memory.                                                                                                                                   |
| <i>value</i>   | The initial value of the semaphore. A positive value (i.e. greater than zero) indicates an unlocked semaphore, and a <i>value</i> of 0 (zero) indicates a locked semaphore. This value must not exceed SEM_VALUE_MAX. |

## Library:

libc

## Description:

The *sem\_init()* function initializes the unnamed semaphore referred to by the *sem* argument. The initial counter value of this semaphore is specified by the *value* argument.

You can use the initialized semaphore in subsequent calls to *sem\_wait()*, *sem\_trywait()*, *sem\_post()*, and *sem\_destroy()*. An initialized semaphore is valid until it's destroyed by the *sem\_destroy()* function, or until the memory where the semaphore resides is released.

If the *pshared* argument is nonzero, then the semaphore can be shared between processes via shared memory. Any process can then use *sem*

with the *sem\_wait()*, *sem\_trywait()*, *sem\_post()* and *sem\_destroy()* functions.



---

Don't mix named semaphore operations (*sem\_open()* and *sem\_close()*) with unnamed semaphore operations (*sem\_init()* and *sem\_destroy()*) on the same semaphore.

---

## Returns:

- 0 Success. The semaphore referred to by *sem* is initialized.
- 1 An error occurred (*errno* is set).

## Errors:

- EAGAIN A resource required to initialize the semaphore has been exhausted.
- EINVAL The *value* argument exceeds SEM\_VALUE\_MAX.
- EPERM The process lacks the appropriate privileges to initialize the semaphore.
- ENOSPC A resource required to initialize the semaphore has been exhausted.
- ENOSYS The *sem\_init()* function isn't supported.

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**Caveats:**

Don't initialize the same semaphore from more than one thread. It's best to set up semaphores before starting any threads.

**See also:**

*errno, sem\_destroy(), sem\_post(), sem\_trywait(), sem\_wait()*

## ***sem\_open()***

© 2004, QNX Software Systems Ltd.

*Create or access a named semaphore*

### **Synopsis:**

```
#include <semaphore.h>

sem_t * sem_open(const char * sem_name,
 int oflags,
 ...);
```

### **Arguments:**

*sem\_name*     The name of the semaphore that you want to create or access; see below.

*oflags*        Flags that affect how the function creates a new semaphore. This argument is a combination of:

- O\_CREAT
- O\_EXCL



---

Don't set *oflags* to O\_RDONLY, O\_RDWR, or O\_WRONLY. A semaphore's behavior is undefined with these flags. The QNX libraries silently ignore these options, but they may reduce your code's portability.

---

For more information, see below.

If you set O\_CREAT in *oflags*, you must also pass the following arguments:

**mode\_t mode**     The semaphore's mode (just like file modes). For portability, you should set the read, write, *and* execute bits to the same value. An easy way of doing this is to use the constants from **<sys/stat.h>**:

- S\_IRWXG for group access.
- S\_IRWXO for other's access.
- S\_IRWXU for your own access.



For more information, see “Access permissions” in the documentation for *stat()*.

**unsigned int** *value*

The initial value of the semaphore. A positive value (i.e. greater than zero) indicates an unlocked semaphore, and a *value* of 0 (zero) indicates a locked semaphore. This value must not exceed SEM\_VALUE\_MAX.

## Library:

**libc**

## Description:

The *sem\_open()* function creates or accesses a named semaphore. Named semaphores are slower than the unnamed semaphores created with *sem\_init()*. Semaphores persist as long as the system is up.

The *sem\_open()* function returns a semaphore descriptor that you can use with *sem\_wait()*, *sem\_trywait()*, and *sem\_post()*. You can use it until you call *sem\_close()*.

The *sem\_name* argument is interpreted as follows:

| <i>name</i>            | Pathname space entry      |
|------------------------|---------------------------|
| <i>entry</i>           | <i>CWD/entry</i>          |
| <i>/entry</i>          | <i>/dev/sem/entry</i>     |
| <i>entry/newentry</i>  | <i>CWD/entry/newentry</i> |
| <i>/entry/newentry</i> | <i>/entry/newentry</i>    |

where *CWD* is the current working directory for the program at the point that it calls *mq\_open()*.



---

If you want to create or access a semaphore on another node, you have to specify the name as `/net/node/sem_location`.

---

The *oflags* argument is used only for semaphore creation. When creating a new semaphore, you can set *oflags* to O\_CREAT or (O\_CREAT | O\_EXCL):

O\_CREAT      Create a new named semaphore. If you set this bit, you must provide the *mode* and *value* arguments to *sem\_open()*.

O\_EXCL      When creating a new named semaphore, O\_EXCL causes *sem\_open()* to fail if a semaphore with *sem\_name* already exists. Without O\_EXCL, *sem\_open()* attaches to an existing semaphore or creates a new one if *sem\_name* doesn't exist.



---

Don't mix named semaphore operations (*sem\_open()* and *sem\_close()*) with unnamed semaphore operations (*sem\_init()* and *sem\_destroy()*) on the same semaphore.

---

## Returns:

A pointer to the created or accessed semaphore, or -1 for failure (*errno* is set).

## Errors:

EACCES      Either the named semaphore exists and you don't have permission to access it, or you're trying to create a new semaphore and you don't have permission.

EEXIST      You specified O\_CREAT and O\_EXCL in *oflags*, but the semaphore already exists.

EINVAL      The *sem\_name* argument is invalid or, when creating a semaphore, *value* is greater than SEM.VALUE.MAX.

|              |                                                                                                                                                 |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| EINTR        | The call was interrupted by a signal.                                                                                                           |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                  |
| EMFILE       | The process is using too many files or semaphores.                                                                                              |
| ENFILE       | The system ran out of resources and couldn't open the semaphore.                                                                                |
| ENAMETOOLONG | The <i>sem_name</i> argument is longer than (NAME_MAX - 8).                                                                                     |
| ENOENT       | Either the <b>mqueue</b> manager isn't running, or the <i>sem_name</i> argument doesn't exist and you didn't specify O_CREAT in <i>oflags</i> . |
| ENOSPC       | There's insufficient space to create a new named semaphore.                                                                                     |
| ENOSYS       | The <i>sem_open()</i> function isn't implemented for the filesystem specified in <i>sem_name</i> .                                              |

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **Caveats:**

The `mqueue` manager must be running for applications to use named semaphores.

### **See also:**

*sem\_close()*, *sem\_destroy()*, *sem\_init()*, *sem\_post()*, *sem\_trywait()*,  
*sem\_unlink()*, *sem\_wait()*

`mqueue` in the *Utilities Reference*

## Synopsis:

```
#include <semaphore.h>

int sem_post(sem_t * sem);
```

## Arguments:

*sem* A pointer to the **sem\_t** object for the semaphore whose value you want to increment.

## Library:

libc

## Description:

The *sem\_post()* function increments the semaphore referenced by the *sem* argument. If any processes are currently blocked waiting for the semaphore, then one of these processes will return successfully from its call to *sem\_wait*.

The process to be unblocked is determined in accordance with the scheduling policies in effect for the blocked processes. The highest priority waiting process is unblocked, and if there is more than one highest priority process blocked waiting for the semaphore, then the highest priority process that has been waiting the longest is unblocked.

The *sem\_post()* function is reentrant with respect to signals, and can be called from a signal handler.

## Returns:

0 Success.  
-1 An error occurred (*errno* is set).

## **Errors:**

- EINVAL      Invalid semaphore descriptor *sem*.
- ENOSYS      The *sem\_post()* function isn't supported.

## **Classification:**

POSIX 1003.1

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*errno*, *sem\_destroy()*, *sem\_init()*, *sem\_trywait()*, *sem\_wait()*

### **Synopsis:**

```
#include <semaphore.h>
#include <time.h>

int sem_timedwait(
 sem_t * sem,
 const struct timespec * abs_timeout);
```

### **Arguments:**

|                    |                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| <i>sem</i>         | The semaphore that you want to wait on.                                                                                                |
| <i>abs_timeout</i> | A pointer to a <b>timespec</b> structure that specifies the maximum time to wait to lock the semaphore, expressed as an absolute time. |

### **Library:**

libc

### **Description:**

The *sem\_timedwait()* function locks the semaphore referenced by *sem* as in the *sem\_wait()* function. However, if the semaphore can't be locked without waiting for another process or thread to unlock the semaphore by calling *sem\_post()*, the wait is terminated when the specified timeout expires.

The timeout expires when the absolute time specified by *abs\_timeout* passes, as measured by the clock on which timeouts are based (i.e. when the value of that clock equals or exceeds *abs\_timeout*), or if the absolute time specified by *abs\_timeout* has already been passed at the time of the call. The timeout is based on the CLOCK\_REALTIME clock.

## Returns:

- 0 The calling process successfully performed the semaphore lock operation on the semaphore designated by *sem*.
- 1 The call was unsuccessful (*errno* is set). The state of the semaphore is unchanged.

## Errors:

- EDEADLK A deadlock condition was detected.
- EINTR A signal interrupted this function.
- EINVAL Invalid semaphore *sem*, or the thread would have blocked, and the *abs\_timeout* parameter specified a nanoseconds field value less than zero or greater than or equal to 1000 million.
- ETIMEDOUT The semaphore couldn't be locked before the specified timeout expired.

## Examples:

```
#include <stdio.h>
#include <semaphore.h>
#include <time.h>

main() {

 struct timespec tm;
 sem_t sem;
 int i=0;

 sem_init(&sem, 0, 0);

 do {
 clock_gettime(CLOCK_REALTIME, &tm);
 tm.tv_sec += 1;
 i++;
 printf("i=%d\n",i);
 if (i==10) {
 sem_post(&sem);
 }
 } while (sem_timedwait(&sem, &tm) == -1);
}
```



```
 printf("Semaphore acquired after %d timeouts\n", i);
 return;
 }
```

## Classification:

POSIX 1003.1d (draft)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*sem\_post()*, *sem\_trywait()*, *sem\_wait()*, *time()*, **timespec**

## ***sem\_trywait()***

© 2004, QNX Software Systems Ltd.

*Wait on a semaphore, but don't block*

### **Synopsis:**

```
#include <semaphore.h>

int sem_trywait(sem_t * sem);
```

### **Arguments:**

*sem*     A pointer to the **sem\_t** object for the semaphore that you want to wait on.

### **Library:**

**libc**

### **Description:**

The *sem\_trywait()* function decrements the semaphore if the semaphore's value is greater than zero; otherwise, the function simply returns.

### **Returns:**

0        The semaphore was successfully decremented.  
-1       The state of the semaphore is unchanged (*errno* is set).

### **Errors:**

EAGAIN     The semaphore was already locked, so it couldn't be immediately locked by the *sem\_trywait()* function.  
EDEADLK    A deadlock condition was detected.  
EINVAL     Invalid semaphore descriptor *sem*.  
EINTR      A signal interrupted this function.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sem\_destroy()*, *sem\_init()*, *sem\_post()*, *sem\_wait()*

## ***sem\_unlink()***

© 2004, QNX Software Systems Ltd.

*Destroy a named semaphore*

### **Synopsis:**

```
#include <semaphore.h>

int sem_unlink(const char * sem_name);
```

### **Arguments:**

*sem\_name*      The name of the semaphore that you want to destroy.

### **Library:**

`libc`

### **Description:**

The *sem\_unlink()* function destroys the named semaphore, *sem\_name*. Open semaphores are removed the same way that *unlink()* removes open files; the processes that have the semaphore open can still use it, but the semaphore will disappear as soon as the last process uses *sem\_close()* to close it.

Any attempt to use *sem\_open()* on an unlinked semaphore will refer to a *new* semaphore.

Semaphores are persistent as long as the system remains up.

### **Returns:**

0      Success.

-1     An error occurred (*errno* is set).

### **Errors:**

EACCESS      You don't have permission to unlink the semaphore.

ELOOP        Too many levels of symbolic links or prefixes.

ENOENT       The semaphore *sem\_name* doesn't exist.

ENAMETOOLONG

The *sem\_name* argument is longer than (NAME\_MAX - 8).

ENOSYS

The *sem\_unlink()* function isn't implemented for the filesystem specified in *path*.

### Classification:

POSIX 1003.1

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### Caveats:

The **mqueue** manager must be running for applications to use named semaphores.

### See also:

*sem\_open()*, *sem\_close()*, *sem\_wait()*, *sem\_trywait()*, *sem\_post()*

**mqueue** in the *Utilities Reference*

## ***sem\_wait()***

© 2004, QNX Software Systems Ltd.

*Wait on a semaphore*

### **Synopsis:**

```
#include <semaphore.h>

int sem_wait(sem_t * sem);
```

### **Arguments:**

*sem*     A pointer to the **sem\_t** object for the semaphore that you want to wait on.

### **Library:**

**libc**

### **Description:**

The *sem\_wait()* function decrements the semaphore referred to by the *sem* argument. If the semaphore value is not greater than zero, then the calling process blocks until it can decrement the counter, or the call is interrupted by signal.

Some process should eventually call *sem\_post()* to increment the semaphore.

### **Returns:**

- 0     The semaphore was successfully decremented.
- 1    The state of the semaphore is unchanged (*errno* is set).

### **Errors:**

- EDEADLK     A deadlock condition was detected.
- EINVAL      Invalid semaphore descriptor *sem*.
- EINTR        A signal interrupted this function.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sem\_destroy()*, *sem\_init()*, *sem\_post()*, *sem\_trywait()*

## **send()**

© 2004, QNX Software Systems Ltd.

*Send a message to a connected socket*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t send(int s,
 const void * msg,
 size_t len,
 int flags);
```

### **Arguments:**

- s*           The descriptor for the socket; see *socket()*.
- msg*          A pointer to the message that you want to send.
- len*          The length of the message.
- flags*        A combination of the following:
- MSG\_OOB — process out-of-band data. Use this bit when you send “out-of-band” data on sockets that support this notion (e.g. SOCK\_STREAM). The underlying protocol must also support out-of-band data.
  - MSG\_DONTROUTE — bypass routing; create a direct interface. You normally use this bit only in diagnostic or routing programs.



---

The tiny TCP/IP stack doesn't support MSG\_OOB and MSG\_DONTROUTE. For more information, see `npm-ttcpip.so` in the *Utilities Reference*.

---

### **Library:**

`libsocket`



## Description:

The *send()*, *sendto()*, and *sendmsg()* functions are used to transmit a message to another socket. The *send()* function can be used only when the socket is in a *connected* state, while *sendto()* and *sendmsg()* can be used at any time.

The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, the error EMSGSIZE is returned, and the message isn't transmitted.

No indication of failure to deliver is implicit in a *send()*. Locally detected errors are indicated by a return value of -1.

If no message space is available at the socket to hold the message to be transmitted, then *send()* normally blocks, unless the socket has been placed in nonblocking I/O mode. You can use *select()* to determine when it's possible to send more data.

## Returns:

The number of bytes sent, or -1 if an error occurs (*errno* is set).

## Errors:

|              |                                                                                                            |
|--------------|------------------------------------------------------------------------------------------------------------|
| EBADF        | An invalid descriptor was specified.                                                                       |
| EDESTADDRREQ | A destination address is required.                                                                         |
| EFAULT       | An invalid user space address was specified for a parameter.                                               |
| EMSGSIZE     | The socket requires that the message be sent atomically, but the size of the message made this impossible. |
| ENOBUFS      | The system couldn't allocate an internal buffer. The operation may succeed when buffers become available.  |
| ENOTSOCK     | The argument <i>s</i> isn't a socket.                                                                      |

EPIPE            An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.

EWOULDBLOCK            The socket is marked nonblocking and the requested operation would block.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*getsockopt(), ioctl(), recv(), select(), sendmsg(), sendto(), socket(), write()*

**Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t sendmsg(int s,
 const struct msghdr * msg,
 int flags);
```

**Arguments:**

- s* The descriptor for the socket; see *socket()*.
- msg* A pointer to the message that you want to send. For a description of the **msghdr** structure, see *recvmsg()*.
- flags* A combination of the following:
- MSG\_OOB — process out-of-band data. Use this bit when you send “out-of-band” data on sockets that support this notion (e.g. SOCK\_STREAM). The underlying protocol must also support out-of-band data.
  - MSG\_DONTROUTE — bypass routing; create a direct interface. You normally use this bit only in diagnostic or routing programs.



---

The tiny TCP/IP stack doesn't support MSG\_OOB and MSG\_DONTROUTE. For more information, see **npm-ttcpip.so** in the *Utilities Reference*.

---

**Library:**

**libsocket**

**Description:**

The *sendmsg()* function is used to transmit a message to another socket. You can use *send()* only when the socket is in a *connected* state; you can use *sendmsg()* at any time.

No indication of failure to deliver is implicit in a *sendmsg()*. Locally detected errors are indicated by a return value of -1.

If no message space is available at the socket to hold the message to be transmitted, then *sendmsg()* normally blocks, unless the socket has been placed in nonblocking I/O mode. You can use *select()* to determine when it's possible to send more data.

## Returns:

The number of bytes sent, or -1 if an error occurs (*errno* is set).

## Errors:

|              |                                                                                                            |
|--------------|------------------------------------------------------------------------------------------------------------|
| EBADF        | An invalid descriptor was specified.                                                                       |
| EDESTADDRREQ | A destination address is required.                                                                         |
| EFAULT       | An invalid user space address was specified for a parameter.                                               |
| EMSGSIZE     | The socket requires that the message be sent atomically, but the size of the message made this impossible. |
| ENOBUFS      | The system couldn't allocate an internal buffer. The operation may succeed when buffers become available.  |
| ENOTSOCK     | The argument <i>s</i> isn't a socket.                                                                      |
| EWOULDBLOCK  | The socket is marked nonblocking and the requested operation would block.                                  |

## Classification:

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*getsockopt(), ioctl(), recv(), select(), send(), sendto(), socket(), write()*

## **sendto()**

© 2004, QNX Software Systems Ltd.

*Send a message to a socket at a specific address*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

ssize_t sendto(int s,
 const void * msg,
 size_t len,
 int flags,
 const struct sockaddr * to,
 socklen_t tolen);
```

### **Arguments:**

- s*            The descriptor for the socket; see *socket()*.
- msg*          A pointer to the message that you want to send.
- len*          The length of the message.
- flags*        A combination of the following:
- MSG\_OOB — process out-of-band data. Use this bit when you send “out-of-band” data on sockets that support this notion (e.g. SOCK\_STREAM). The underlying protocol must also support out-of-band data.
  - MSG\_DONTROUTE — bypass routing; create a direct interface. You normally use this bit only in diagnostic or routing programs.



---

The tiny TCP/IP stack doesn't support MSG\_OOB and MSG\_DONTROUTE. For more information, see `npm-ttcpip.so` in the *Utilities Reference*.

---

*to*            A pointer to a `sockaddr` object that specifies the address of the target.

*tolen*        A `socklen_t` object that specifies the size of the *to* address.

**Library:**`libsocket`**Description:**

The *sendto()* function is used to transmit a message to another socket. You can use *send()* only when the socket is in a *connected* state; you can use *sendto()* at any time.

The address of the target is given by *to*, with *tolen* specifying its size. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, the error EMSGSIZE is returned, and the message isn't transmitted.

No indication of failure to deliver is implicit in a *sendto()*. Locally detected errors are indicated by a return value of -1.

If no message space is available at the socket to hold the message to be transmitted, then *sendto()* normally blocks, unless the socket has been placed in nonblocking I/O mode. You can use *select()* to determine when it's possible to send more data.

**Returns:**

The number of bytes sent, or -1 if an error occurs (*errno* is set).

**Errors:**

|              |                                                                                                            |
|--------------|------------------------------------------------------------------------------------------------------------|
| EBADF        | An invalid descriptor was specified.                                                                       |
| EDESTADDRREQ | A destination address is required.                                                                         |
| EFAULT       | An invalid user space address was specified for a parameter.                                               |
| EMSGSIZE     | The socket requires that the message be sent atomically, but the size of the message made this impossible. |

- ENOBUFS      The system couldn't allocate an internal buffer. The operation may succeed when buffers become available.
- ENOTSOCK    The argument *s* isn't a socket.
- EWOULDBLOCK  
              The socket is marked nonblocking and the requested operation would block.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*getsockopt(), ioctl(), recv(), select(), send(), sendmsg(), socket(), write()*



**Synopsis:**

```
#include <netdb.h>

struct servent {
 char * s_name;
 char ** s_aliases;
 int s_port;
 char * s_proto;
};
```

**Description:**

This structure is used to hold the broken-out fields of a line in the network services database, */etc/services*. The members of this structure are:

|                  |                                                                                               |
|------------------|-----------------------------------------------------------------------------------------------|
| <i>s_name</i>    | The name of the service.                                                                      |
| <i>s_aliases</i> | A zero-terminated list of alternate names for the service.                                    |
| <i>s_port</i>    | The port number that the service resides at. Port numbers are returned in network byte order. |
| <i>s_proto</i>   | The name of the protocol to use when contacting the service.                                  |

**Classification:**

Unix, POSIX 1003.1-2001

**See also:**

*endservent()*, *getservbyname()*, *getservbyport()*, *getservent()*,  
*setservent()*

*/etc/services* in the *Utilities Reference*

## ***setbuf()***

© 2004, QNX Software Systems Ltd.

*Associate a buffer with a stream*

### **Synopsis:**

```
#include <stdio.h>

void setbuf(FILE *fp,
 char *buffer);
```

### **Arguments:**

*fp*            The stream that you want to associate with a buffer.

*buffer*        NULL, or a pointer to the buffer; see below.

### **Library:**

libc

### **Description:**

The *setbuf()* function associates the supplied *buffer* with the stream specified by *fp*. If you want to call *setbuf()*, you must call it after opening the stream, but before doing any reading, writing, or seeking.

If *buffer* is NULL, all input/output for the stream is completely unbuffered. If *buffer* isn't NULL, then it must point to an array that's at least BUFSIZ characters long, and all input/output is fully buffered.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *buffer;
 FILE *fp;

 buffer = (char *)malloc(BUFSIZ);
 if(buffer == NULL) {
 return EXIT_FAILURE;
 }

 fp = fopen("some_file", "r");
```

```
 setbuf(fp, buffer);

 /* . */
 /* . */
 /* . */

 fclose(fp);

 free(buffer);

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*fopen()*, *setvbuf()*

## **setbuffer()**

© 2004, QNX Software Systems Ltd.

*Assign block buffering to a stream*

### **Synopsis:**

```
#include <unix.h>

void setbuffer(FILE *iop,
 char *abuf,
 size_t asize);
```

### **Arguments:**

*iop*        The stream that you want to set the buffering for.

*abuf*       NULL, or a pointer to the buffer that you want the stream to use.

*asize*      The size of the buffer.

### **Library:**

libc

### **Description:**

The *setbuffer()* and *setlinebuf()* functions assign buffering to a stream. The types of buffering available are:

|                |                                                                                                 |
|----------------|-------------------------------------------------------------------------------------------------|
| Unbuffered     | Information appears on the destination file or terminal as soon as written.                     |
| Block-buffered | Many characters are saved and written as a block.                                               |
| Line-buffered  | Characters are saved until either a newline is encountered or input is read from <i>stdin</i> . |

You can use *fflush()* to force the block out early. Normally all files are block-buffered. A buffer is obtained from *malloc()* when you perform the first *getc()* or *putc()* on the file. If the standard stream *stdout* refers to a terminal, it's line-buffered. The standard stream *stderr* is unbuffered by default.

If you want to use *setbuffer()*, you must call it after opening the stream, but before doing any reading or writing. It uses the character array *abuf*, whose size is given by *asize*, instead of an automatically allocated buffer. If *abuf* is NULL, input and output are completely unbuffered. A manifest constant BUFSIZ, defined in the `<stdio.h>` header, tells how large an array is needed:

```
char buf[BUFSIZ];
```

You can use *freopen()* to change a stream from unbuffered or line-buffered to block buffered. To change a stream from block-buffered or line-buffered to unbuffered, call *freopen()*, and then call *setbuf()* with a buffer argument of NULL.

## Classification:

Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

A common source of error is allocating buffer space as an **automatic** variable in a code block, and then failing to close the stream in the same block.

## See also:

*fclose()*, *fflush()*, *fopen()*, *fread()*, *freopen()*, *getc()*, *malloc()*, *printf()*, *putc()*, *puts()*, *setbuf()*, *setlinebuf()*, *setvbuf()*

## ***setdomainname()***

© 2004, QNX Software Systems Ltd.

*Set the domain name of the current host*

---

### **Synopsis:**

```
#include <unistd.h>

int setdomainname(const char * name,
 size_t namelen);
```

### **Arguments:**

*name*            The domain name.

*namelen*        The length of the name.

### **Library:**

`libsocket`

### **Description:**

The *setdomainname()* function sets the domain *name* of the host machine. Only the superuser (`root`) can use this function and even then, the function is normally used only when bootstrapping a system.

### **Returns:**

0            Success.

-1          Failure; *errno* is set.

### **Errors:**

EFAULT      The *name* or *namelen* parameters gave an invalid address.

EPERM        The caller tried to set the domain name without being the superuser.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*getdomainname()*

## ***setegid()***

© 2004, QNX Software Systems Ltd.

*Set the effective group ID for a process*

---

### **Synopsis:**

```
#include <unistd.h>

int setegid(gid_t gid);
```

### **Arguments:**

*gid*     The effective group ID that you want to use for the process.

### **Library:**

`libc`

### **Description:**

The *setegid()* function lets the calling process set the effective group ID based on the following:

- If the process is the superuser, the *setegid()* function sets the effective group ID to *gid*.
- If the process isn't the superuser, but *gid* is equal to the real group ID or saved set-group ID, *setegid()* sets the effective group ID to *gid*.

The real and saved group ID aren't changed.



If a set-group ID process sets its effective group ID to its real group ID, it can still set its effective group ID back to the saved set-group ID.

---

The "superuser" is defined as any process with an effective user ID of 0, or an effective user ID of `root`.



**Returns:**

Zero for success, or -1 if an error occurs (*errno* is set).

**Errors:**

EINVAL     The value of *gid* is out of range.

EPERM     The process isn't the superuser, and *gid* doesn't match the real group ID or the saved set-group ID.

**Examples:**

```
/*
 * This process sets its effective group ID to 2
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <unistd.h>

int main(void)
{
 gid_t oegid;

 oegid = getegid();
 if(setegid(2) == -1) {
 perror("setegid");
 return EXIT_FAILURE;
 }

 printf("Was effective group %d, is 2\n", oegid);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1a

**Safety**

Cancellation point   No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*errno, getegid(), seteuid(), setgid(), setuid()*

## Synopsis:

```
#include <stdlib.h>

int setenv(const char* name,
 const char* value,
 int overwrite);
```

## Arguments:

|                  |                                                                                                                                    |
|------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <i>name</i>      | The name of the environment variable that you want to set.                                                                         |
| <i>value</i>     | NULL, or the value for the environment variable; see below.                                                                        |
| <i>overwrite</i> | A nonzero value if you want the function to overwrite the variable if it exists, or 0 if you don't want to overwrite the variable. |

## Library:

libc

## Description:

The *setenv()* function sets the environment variable *name* to *value*. If *name* doesn't exist in the environment, it's created; if *name* exists and *overwrite* is nonzero, the variable's old value is overwritten with *value*; otherwise, it isn't changed.

Copies of the specified *name* and *value* are placed in the environment.

If *value* is NULL, the environment variable specified by *name* is removed from the environment.



---

The value of the global *environ* pointer could be changed by a call to the *setenv()* function.

---

Environment variable names are case-sensitive.

## Returns:

0           Success.  
Nonzero     An error occurred (*errno* is set).

## Errors:

ENOMEM     Not enough memory to allocate a new environment variable.

## Examples:

Change the string assigned to **INCLUDE** and then display the new string:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char* path;

 if(setenv("INCLUDE",
 "/usr/nto/include:/home/fred/include",
 1) == 0) {
 if((path = getenv("INCLUDE")) != NULL) {
 printf("INCLUDE=%s\n", path);
 }
 }

 return EXIT_SUCCESS;
}
```

## **Classification:**

POSIX 1003.1a

### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## **Caveats:**

The *setenv()* function manipulates the environment pointed to by the global *environ* variable.

## **See also:**

*clearenv()*, *errno*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *getenv()*, *putenv()*, *searchenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *system()*, *unsetenv()*

## **seteuid()**

© 2004, QNX Software Systems Ltd.

*Set the effective user ID*

### **Synopsis:**

```
#include <unistd.h>

int seteuid(uid_t uid);
```

### **Arguments:**

*uid*     The effective user ID that you want to use for the process.

### **Library:**

`libc`

### **Description:**

The *seteuid()* function lets the calling process set the effective user ID, based on the following:

- If the process is the superuser, the *seteuid()* function sets the effective user ID to *uid*.
- If the process isn't the superuser, and *uid* is equal to the real user ID or saved set-user ID, *seteuid()* sets the effective user ID to *uid*.

The real and saved user IDs aren't changed.



---

If a set-UID process sets its effective user ID to its real user ID, it can still set its effective user ID back to the saved set-UID.

---

The "superuser" is defined as any process with an effective user ID of 0, or an effective user ID of `root`.

### **Returns:**

- 0     Success.
- 1    An error occurred (*errno* is set).

**Errors:**

- EINVAL     The value of *uid* is out of range.
- EPERM     The process isn't the superuser, and *uid* doesn't match the real user ID or the saved set-user ID.

**Examples:**

```
/*
 * This process sets its effective userid to 0 (root).
 */

#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 uid_t oeuid;

 oeuid = geteuid();
 if(seteuid(0) == -1) {
 perror("seteuid");
 return EXIT_FAILURE;
 }

 printf("effective userid now 0, was %d\n",
 oeuid);

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1a

**Safety**

Cancellation point   No

Interrupt handler    No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*errno, geteuid(), setegid(), setuid(), setgid()*



### **Synopsis:**

```
#include <unistd.h>

int setgid(gid_t gid);
```

### **Arguments:**

*gid*      The group ID that you want to use for the process.

### **Library:**

`libc`

### **Description:**

The *setgid()* function lets the calling process set the real, effective and saved group IDs, based on the following:

- If the process is the superuser, the *setgid()* function sets the real group ID, effective group ID and saved group ID to *gid*.
- If the process isn't the superuser, but *gid* is equal to the real group ID, *setgid()* sets the effective group ID to *gid*; the real and saved group IDs aren't changed.

This function doesn't change any supplementary group IDs of the calling process.

If you wish to change only the effective group ID, and even if you are the superuser, you should consider using the *setegid()* function.

The "superuser" is defined as any process with an effective user ID of 0, or an effective user ID of `root`.

### **Returns:**

- 0      Success.
- 1     An error occurred; *errno* is set to indicate the error.

## Errors:

- EINVAL The value of *gid* is invalid.
- EPERM The process doesn't have appropriate privileges, and *gid* doesn't match the real group ID.

## Examples:

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 gid_t ogid;

 ogid = getgid();
 if(setgid(2) == -1) {
 perror("setgid");
 return EXIT_FAILURE;
 }
 printf("group id is now 2, was %d\n", ogid);
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno, setegid(), seteuid(), setuid()*

## ***setgrent()***

© 2004, QNX Software Systems Ltd.

*Rewind to the start of the group database file*

### **Synopsis:**

```
#include <grp.h>

int setgrent(void);
```

### **Library:**

libc

### **Description:**

The *setgrent()* function rewinds to the start of the group name database file. It's provided for programs that make multiple lookups in the group database (using the *getgrgid()* and *getgrnam()* calls) to avoid the default opening and closing of the group database for each access.

### **Returns:**

0      Success.  
-1     An error occurred.

### **Errors:**

The *setgrent()* function uses *fopen()*. As a result, *errno* can be set to an error for the *fopen()* call.

### **Classification:**

Standard Unix

#### **Safety**

---

Cancellation point    Yes  
Interrupt handler      No

*continued...*

**Safety**

---

|                |    |
|----------------|----|
| Signal handler | No |
| Thread         | No |

**See also:**

*endgrent()*, *getgrent()*

# **setgroups()**

© 2004, QNX Software Systems Ltd.

*Set supplementary group IDs*

## **Synopsis:**

```
#include <unistd.h>

int setgroups(int ngroups,
 const gid_t *gidset);
```

## **Arguments:**

*ngroups*     The number of entries in the *gidset* array.

*gidset*     An array of the supplementary group IDs that you want to assign to the current user. This number of entries in this array can't exceed NGROUPS\_MAX.

## **Library:**

libc

## **Description:**

The *setgroups()* function sets the group access list of the current user to the array of group IDs in *gidset*.



---

Only **root** can set new groups.

---

## **Returns:**

0, or -1 if an error occurred (*errno* is set).

## **Errors:**

EFAULT     The *gidset* argument isn't a valid pointer.

EPERM     The caller isn't **root**.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*getgroups()*, *initgroups()*

## ***sethostent()***

© 2004, QNX Software Systems Ltd.

*Open the host database file*

### **Synopsis:**

```
#include <netdb.h>

void sethostent(int stayopen);
```

### **Arguments:**

*stayopen*      Nonzero if you want all queries to the name server to use TCP and you want the connection to be retained after each call to *gethostbyname()* or *gethostbyaddr()*.  
If the *stayopen* flag is zero, queries use UDP datagrams.

### **Library:**

`libsocket`

### **Description:**

The *sethostent()* routine opens the host database file.

You can use the *sethostent()* function to request the use of a connected TCP socket for queries.

### **Classification:**

Standard Unix, POSIX 1003.1-2001

#### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |



## **Caveats:**

This function uses static data; if the data is needed for future use, it should be copied before any subsequent calls overwrite it.

## **See also:**

*endhostent()*, *gethostbyaddr()*, *gethostbyname()*, *gethostent()*,  
*gethostent\_r()*, **hostent**

*/etc/hosts*, */etc/resolv.conf* in the *Utilities Reference*

# ***sethostname()***

© 2004, QNX Software Systems Ltd.

*Set the name of the current host*

## **Synopsis:**

```
#include <unistd.h>

int sethostname(const char * name,
 size_t namelen);
```

## **Arguments:**

*name*            The name that you want to use for the host machine. Hostnames are limited to MAXHOSTNAMELEN characters (defined in <sys/param.h>).

*namelen*        The length of the name.

## **Library:**

libc

## **Description:**

The *sethostname()* function sets the name of the host machine to be *name*. Only the superuser can call this function; this is normally done only at boot time.



---

This function sets the value of the **\_CS\_HOSTNAME** configuration string, not that of the **HOSTNAME** environment variable.

---

## **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

## **Errors:**

EFAULT      Either *name* or *namelen* gave an invalid address.

EPERM        Although the caller wasn't the superuser, it tried to set the hostname.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

This function is restricted to the superuser, and is normally used only at boot time.

## See also:

*gethostname()*

# SETIOV()

© 2004, QNX Software Systems Ltd.

Fill in the fields of an `iov_t` structure

## Synopsis:

```
#include <unistd.h>

void SETIOV(iov_t *msg,
 void *addr,
 size_t len);
```

## Arguments:

*msg* A pointer to the `iov_t` structure structure that you want to set.

*addr* The value you want to use for the structure's `iov_base` member.

*len* The value you want to use for the structure's `iov_len` member.

## Description:

The `SETIOV()` macro fills in the fields of an `iov_t` message structure. The `iov_t` structure consists of two fields:

```
typedef struct iovec {
 void *iov_base;
 size_t iov_len;
} iov_t;
```



---

`SETIOV()` doesn't make a copy of the data that `addr` points to; it just copies the pointer.

---

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*GETIOVBASE(), GETIOVLEN(), MsgKeyData(), MsgReady(),  
MsgReceivev(), MsgReplyv(), MsgSendv(), MsgWritev()*

## **setitimer()**

© 2004, QNX Software Systems Ltd.

*Set the value of an interval timer*

### **Synopsis:**

```
#include <sys/time.h>

int setitimer (int which,
 const struct itimerval *value,
 struct itimerval *ovalue);
```

### **Arguments:**

*which*      The interval time whose value you want to set. Currently, this must be ITIMER\_REAL.

*value*      A pointer to a `itimerval` structure that specifies the value that you want to set the interval timer to.

*ovalue*      NULL, or a pointer to a `itimerval` structure where the function can store the old value of the interval timer.

### **Library:**

`libc`

### **Description:**

The system provides each process with interval timers, defined in `<sys/time.h>`. The `setitimer()` function sets the value of the timer specified by *which* to the value specified in the structure pointed to by *value*, and if *ovalue* isn't NULL, stores the previous value of the timer in the structure it points to.

A timer value is defined by the `itimerval` structure (see `gettimeofday()` for the definition of `timeval`), which includes the following members:

```
struct timeval it_interval; /* timer interval */
struct timeval it_value; /* current value */
```

The *it\_value* member indicates the time to the next timer expiration. The *it\_interval* member specifies a value to be used in reloading

*it\_value* when the timer expires. Setting *it\_value* to 0 disables a timer, regardless of the value of *it\_interval*. Setting *it\_interval* to 0 disables a timer after its next expiration (assuming *it\_value* is nonzero).

Time values smaller than the resolution of the system clock are rounded up to the resolution of the system clock.

The only supported timer is ITIMER\_REAL, which decrements in real time. A SIGALRM signal is delivered when this timer expires.

The SIGALRM so generated isn't maskable on this bound thread by any signal-masking function, *pthread\_sigmask()*, or *sigprocmask()*.

## Returns:

- 0 Success.
- 1 An error occurred; *errno* is set.

## Errors:

- EINVAL The specified number of seconds is greater than 100,000,000, the number of microseconds is greater than or equal to 1,000,000, or the *which* argument is unrecognized.

## Classification:

Standard Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

All flags to *setitimer()* other than `ITIMER_REAL` behave as documented only with “bound” threads. Their ability to mask the signal works only with bound threads. If the call is made using one of these flags from an unbound thread, the system call returns -1 and sets *errno* to `EACCES`.

These behaviors are the same for bound or unbound POSIX threads. A POSIX thread with system-wide scope, created by the call:

```
pthread_attr_setscope(&attr, PTHREAD_SCOPE_SYSTEM);
```

is equivalent to a Solaris bound thread. A POSIX thread with local process scope, created by the call:

```
pthread_attr_setscope(&attr, PTHREAD_SCOPE_PROCESS);
```

is equivalent to a Solaris unbound thread.

The microseconds field shouldn't be equal to or greater than one second.

The *setitimer()* function is independent of *alarm()*.

Don't use *setitimer(ITIMER\_REAL)* with the *sleep()* routine. A *sleep()* call wipes out knowledge of the user signal handler for `SIGALRM`.

The granularity of the resolution of the alarm time is platform-dependent.

**See also:**

*alarm()*, *getitimer()*, *gettimeofday()*, *pthread\_attr\_setscope()*, *pthread\_sigmask()*, *sigprocmask()*, *sleep()*, *sysconf()*



### **Synopsis:**

```
#include <setjmp.h>

int setjmp(jmp_buf env);
```

### **Arguments:**

*env*     A buffer where the function can save the calling environment.

### **Library:**

`libc`

### **Description:**

The *setjmp()* function saves the calling environment in its *env* argument for use by the *longjmp()* function.

Error handling can be implemented by using *setjmp()* to record the point to return to following an error. When an error is detected in a function, that function uses *longjmp()* to jump back to the recorded position. The original function that called *setjmp()* must still be active (that is, it can't have returned to the function that called it).

Be careful to ensure that any resources (allocated memory, opened files, etc) are cleaned up properly.



---

**WARNING: Do not use *longjmp()* or *siglongjmp()* to restore an environment saved by a call to *setjmp()* or *sigsetjmp()* in another thread. If you're lucky, your application will crash; if not, it'll look as if it works for a while, until random scribbling on the stack causes it to crash.**

---

### **Returns:**

Zero on the first call, or nonzero if the return is the result of a call to the *longjmp()* function.

## Examples:

```
#include <stdio.h>
#include <setjmp.h>
#include <stdlib.h>

jmp_buf env;

void rtn(void)
{
 printf("about to longjmp()\n");
 longjmp(env, 14);
}

int main(void)
{
 int ret_val;

 ret_val = setjmp(env);

 if(ret_val == 0) {
 printf("after setjmp(): %d\n", ret_val);
 rtn();
 printf("back from rtn(): %d\n", ret_val);
 } else {
 printf("back from longjmp(): %d\n", ret_val);
 }

 return EXIT_SUCCESS;
}
```

produces the output:

```
after setjmp(): 0
about to longjmp()
back from longjmp(): 14
```

## Classification:

ANSI, POSIX 1003.1

### Safety

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*longjmp()*

# **setkey()**

© 2004, QNX Software Systems Ltd.

*Set the key used in encryption*

## **Synopsis:**

```
#include <stdlib.h>

void setkey(const char *__key);
```

## **Arguments:**

*\_\_key*     A 64-character array of binary values (numeric 0 or 1).

## **Library:**

`libc`

## **Description:**

The *setkey()* function allows limited access to the NBS Data Encryption Standard (DES) algorithm itself. It derives a 56-bit key from the given *\_\_key* by dividing the array into groups of 8 and ignoring the last bit in each group.

## **Classification:**

Standard Unix

### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*crypt(), encrypt()*

## ***setlinebuf()***

© 2004, QNX Software Systems Ltd.

*Assign line buffering to a stream*

### **Synopsis:**

```
#include <unix.h>

int setlinebuf(FILE *iop);
```

### **Arguments:**

*iop*     The stream that you want to use line buffering.

### **Library:**

`libc`

### **Description:**

The *setbuffer()* and *setlinebuf()* functions assign buffering to a stream. The types of buffering available are:

|                |                                                                                                 |
|----------------|-------------------------------------------------------------------------------------------------|
| Unbuffered     | Information appears on the destination file or terminal as soon as written.                     |
| Block-buffered | Many characters are saved and written as a block.                                               |
| Line-buffered  | Characters are saved until either a newline is encountered or input is read from <i>stdin</i> . |

You can use *fflush()* to force the block out early. Normally all files are block-buffered. A buffer is obtained from *malloc()* when the first *getc()* or *putc()* is performed on the file. If the standard stream *stdout* refers to a terminal, it's line-buffered. The standard stream *stderr* is unbuffered by default.

You can use *setlinebuf()* to change the buffering on a stream from block-buffered or unbuffered to line-buffered. Unlike *setbuffer()*, you can call *setlinebuf()* at any time that the stream *iop* is active.

You can use *freopen()* to change a stream from unbuffered or line-buffered to block buffered. To change a stream from block-buffered or line-buffered to unbuffered, call *freopen()*, and then call *setbuf()* with a buffer argument of NULL.

**Returns:**

No useful value.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fclose(), fflush(), fopen(), fread(), freopen(), getc(), malloc(), printf(), putc(), puts(), setbuf(), setbuffer(), setvbuf()*

## **setlocale()**

© 2004, QNX Software Systems Ltd.

*Select a program's locale*

### **Synopsis:**

```
#include <locale.h>

char * setlocale(int category,
 const char * locale);
```

### **Arguments:**

*category*      The part of the environment that you want to set; one of:

- LC\_ALL — select the entire locale environment.
- LC\_COLLATE — select only the collating sequence.
- LC\_CTYPE — select only the character-handling information.
- LC\_MESSAGES — specify the language to be used for messages.
- LC\_MONETARY — select only monetary formatting information.
- LC\_NUMERIC — select only the numeric-format environment.
- LC\_TIME — select only the time-related environment.

*locale*      The locale that you want to use. The following built-in locales are offered:

- C (default)
- C-TRADITIONAL
- POSIX

### **Library:**

libc



## Description:

The *setlocale()* function selects a program's *locale*, according to the specified *category* and the specified *locale*.

A locale affects several things:

- The collating sequence (the order in which characters compare with one another) used by *strcoll()* or *wcscoll()*.
- The way certain character-handling functions (such as *isalnum()* and *isalpha()*) operate. The wide-character versions include *iswalnum()* and *iswalpha()*.
- The decimal-point character used in formatted input/output and string conversion (*printf()*, *scanf()*, and friends).
- The format and names used in the string produced by the *strtime()* and *wcsftime()* functions.

See the *localeconv()* function for more information about the locale.

At the start of a program, the default C locale is initialized as if the following call to *setlocale()* appeared at the start of *main()*:

```
(void)setlocale(LC_ALL, "C");
```

## Returns:

A pointer to a system-generated string indicating the previous locale, or NULL if an error occurs.

## Classification:

ANSI

### Safety

Cancellation point No

*continued...*

## **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

## **See also:**

*isalpha()*, *isascii()*, *localeconv()*, *printf()*, *scanf()*, *strcat()*, *strchr()*,  
*strcmp()*, *strcoll()*, *strcpy()*, *strftime()*, *strlen()*, *strpbrk()*, *strspn()*,  
*strtod()*, *strtok()*, *strxfrm()* **tm**,

### **Synopsis:**

```
#include <syslog.h>

int setlogmask(int maskpri);
```

### **Arguments:**

*maskpri*     The new log priority mask; see below.

### **Library:**

`libc`

### **Description:**

The *setlogmask()* function sets the log priority mask to *maskpri* and returns the previous mask. Calls to *syslog()* or *vsyslog()* with a priority that isn't set in *maskpri* are rejected.

You can calculate the mask for an individual priority *pri* with the macro:

```
LOG_MASK (pri) ;
```

You can get the mask for all priorities up to and including *toppri* with the macro:

```
LOG_UPTO (toppri) ;
```

The default allows all priorities to be logged. See the *syslog()* function for a list of the priorities.

### **Returns:**

The previous log mask level.

**Examples:**

See *syslog()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*closelog()*, *openlog()*, *syslog()*, *vsyslog()*,  
**logger**, **syslogd** in the *Utilities Reference*

**Synopsis:**

```
#include <netdb.h>

void setnetent(int stayopen);
```

**Arguments:**

*stayopen*      Nonzero if you don't want the network database to be closed after each call to *getnetbyname()* or *getnetbyaddr()*.

**Library:**

`libsocket`

**Description:**

The *setnetent()* function opens and rewinds the network name database file.

**Files:**

`/etc/networks`  
Network name database file.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**Caveats:**

This function uses static data; if you need the data for future use, you should copy it before any subsequent calls overwrite it.

**See also:**

*endnetent()*, *getnetbyaddr()*, *getnetbyname()*, *getnetent()*, **netent**  
*/etc/networks* in the *Utilities Reference*

## Synopsis:

```
#include <process.h>

int setpgid(pid_t pid,
 pid_t pgid);
```

## Arguments:

*pid*      0, or the ID of the process whose process group you want to set.

*pgid*     0, or the process group ID that you want to join or create.

## Library:

`libc`

## Description:

The *setpgid()* function is used either to join an existing process group or to create a new process group within the session of the calling process. The process group ID of a session leader doesn't change.

The following definitions are worth mentioning:

|               |                                                                                                                  |
|---------------|------------------------------------------------------------------------------------------------------------------|
| Process       | An executing instance of a program, identified by a nonnegative integer called a process ID.                     |
| Process group | A collection of one or more processes, with a unique process group ID. A process group ID is a positive integer. |

On successful completion, the process group ID of the process with a process ID matching *pid* is set to *pgid*. As a special case, you can specify either *pid* or *pgid* as zero, meaning that the process ID of the calling process is to be used.

**Returns:**

- 0 Success.
- 1 An error occurred; *errno* is set.

**Errors:**

- EACCES The value of the *pid* argument matches the process ID of a child process of the calling process, and the child process has successfully executed one of the *exec\**(*)* functions.
- EINVAL The value of *pgid* is less than zero.
- ENOSYS The *setpgid()* function isn't supported by this implementation (included for POSIX compatibility).
- EPERM The calling process doesn't have sufficient privilege to set the process group id *pgid* on process *pid*.
- ESRCH The process *pid* doesn't exist.

**Examples:**

```
/*
 * The process joins process group 0.
 */
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
#include <stdlib.h>
#include <process.h>

int main(void)
{
 if(setpgid(getpid(), 0) == -1) {
 perror("setpgid");
 }
 printf("%d belongs to process group %d\n",
 getpid(), getpgrp());
 return EXIT_SUCCESS;
}
```



**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno, getpgid(), getsid(), setsid()*

# **setpgrp()**

© 2004, QNX Software Systems Ltd.

*Set the process group*

## **Synopsis:**

```
#include <unistd.h>

pid_t setpgrp(void);
```

## **Library:**

libc

## **Description:**

If the calling process isn't already a session leader, *setpgrp()* makes it one by setting its process group ID and session ID to the value of its process ID, and releases its controlling terminal.

## **Returns:**

The new process group ID.

## **Classification:**

Standard Unix

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), fork(), getpid(), getpgrp(), getsid(), kill(), signal()*

## Synopsis:

```
#include <sched.h>

int setprio(pid_t pid,
 int prio);
```

## Arguments:

*pid*      The process ID of the process whose priority you want to set.

*prio*     The new priority.

## Library:

`libc`

## Description:

The *setprio()* function changes the priority of thread 1 of process *pid* to priority *prio*. If *pid* is zero, the priority of the calling thread is set.

## Returns:

The previous priority, or -1 if an error occurred (*errno* is set).

## Errors:

EINVAL     The priority *prio* isn't a valid priority.

EPERM     The calling process doesn't have sufficient privilege to set the priority.

ESRCH     The process *pid* doesn't exist.

## Classification:

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The *getprio()* and *setprio()* functions are included in the QNX Neutrino libraries for porting QNX 4 applications. For new programs, use *sched\_setparam()* or *pthread\_setschedparam()*.

**See also:**

*errno*, *getprio()*, *pthread\_getschedparam()*, *pthread\_setschedparam()*, *sched\_getparam()*, *sched\_get\_priority\_max()*, *sched\_get\_priority\_min()*, *sched\_getscheduler()*, *sched\_setscheduler()*, *sched\_yield()*

**Synopsis:**

```
#include <netdb.h>

void setprotoent(int stayopen);
```

**Arguments:**

*stayopen*      Nonzero if you don't want the database to be closed after each call to *getprotobyname()* or *getprotobynumber()*.

**Library:**

`libsocket`

**Description:**

The *setprotoent()* function opens and rewinds the protocol name database file. If the *stayopen* flag is nonzero, the protocol name database isn't closed after each call to *getprotobyname()* or *getprotobynumber()*.

**Files:**

`/etc/protocols`  
Protocol name database file.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |

*continued...*

**Safety**

---

Thread                      No

**Caveats:**

This function uses static data; if you need the data for future use, you should copy it before any subsequent calls overwrite it.

Currently, only the Internet protocols are understood.

**See also:**

*endprotoent()*, *getprotobyname()*, *getprotobynumber()*, *getprotoent()*,  
**protoent**

*/etc/protocols* in the *Utilities Reference*

**Synopsis:**

```
#include <sys/types.h>
#include <pwd.h>

int setpwent(void);
```

**Library:**

libc

**Description:**

The *setpwent()* function rewinds to the start of the password name database file. It's provided for programs that make multiple lookups in the password database (using the *getpwuid()* and *getpwnam()* calls) to avoid opening and closing the password database for each access.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**See also:**

*endpwent()*, *getpwent()*

## **setregid()**

© 2004, QNX Software Systems Ltd.

*Set real and effective group IDs*

### **Synopsis:**

```
#include <unistd.h>

int setregid(gid_t rgid,
 gid_t egid);
```

### **Arguments:**

*rgid*     The real group ID that you want to use for the process, or -1 if you don't want to change it.

*egid*     The effective group ID that you want to use for the process, or -1 if you don't want to change it.

### **Library:**

libc

### **Description:**

The *setregid()* function sets the real and effective group IDs of the calling process. If *rgid* or *egid* is -1, the corresponding real or effective group ID is left unchanged.

If the effective user ID of the calling process is the superuser, you can set the real group ID and the effective group ID to any legal value.

If the effective user ID of the calling process isn't the superuser, you can set either the real group ID to the saved set-group ID, or the effective group ID to either the saved set-group ID or the real group ID.



If a set-group ID process sets its effective group ID to its real group ID, it can still set its effective group ID back to the saved set-group ID.

---

The "superuser" is defined as any process with an effective user ID of 0, or an effective user ID of `root`.



In either case, if you're changing the real group ID (i.e. *rgid* isn't -1), or you're changing the effective group ID to a value that isn't equal to the real group ID, the saved set-group ID is set to the new effective group ID.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EINVAL The *rgid* or *egid* is out of range.
- EPERM The calling process isn't the superuser, and you tried to change the effective group ID to a value other than the real or saved set-group ID.

Or:

The calling process isn't the superuser, and you tried to change the real group ID to a value other than the effective group ID.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*execve(), getgid(), setreuid(), setuid()*

## Synopsis:

```
#include <unistd.h>

int setreuid(uid_t ruid,
 uid_t euid);
```

## Arguments:

- ruid*     The real user ID that you want to use for the process, or -1 if you don't want to change it.
- euid*     The effective user ID that you want to use for the process, or -1 if you don't want to change it.

## Library:

`libc`

## Description:

The *setreuid()* function sets the real and effective user IDs of the calling process. If *ruid* or *euid* is -1, the corresponding real or effective user ID isn't changed.

If the effective user ID of the calling process is the superuser, you can set the real user ID and the effective user ID to any legal value.

If the effective user ID of the calling process isn't the superuser, you can set either the real user ID to the effective user ID, or the effective user ID to the saved set-user ID or the real user ID.



---

If a set-UID process sets its effective user ID to its real user ID, it can still set its effective user ID back to the saved set-UID.

---

In either case, if you're changing the real user ID (i.e. *ruid* is not -1), or you're changing the effective user ID to a value that isn't equal to the real user ID, the saved set-user ID is set equal to the new effective user ID.

The “superuser” is defined as any process with an effective user ID of 0, or an effective user ID of `root`.

**Returns:**

Zero on success, or -1 if an error occurs (*errno* is set).

**Errors:**

EINVAL     The *ruid* or *euid* is out of range.

EPERM     The calling process isn't the superuser, and you tried to change the effective user ID to a value other than the real or saved set-user ID.

Or

The calling process isn't the superuser, and you tried to change the real user ID to a value other than the effective user ID.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*execve()*, *getuid()*, *setregid()*, *setuid()*

**Synopsis:**

```
#include <sys/resource.h>

int setrlimit(int resource,
 const struct rlimit * rlp);

int setrlimit64(int resource,
 const struct rlimit64 * rlp);
```

**Arguments:**

*resource* The resource whose limit you want to get. For a list of the possible resources, their descriptions, and the actions taken when the current limit is exceeded, see below.

*rlp* A pointer to a `rlimit` or `rlimit64` structure that specifies the limit that you want to set for the resource. The `rlimit` and `rlimit64` structures include at least the following members:

```
rlim_t rlim_cur; /* current (soft) limit */
rlim_t rlim_max; /* hard limit */
```

The `rlim_t` type is an arithmetic data type to which you can cast objects of type `int`, `size_t`, and `off_t` without loss of information.

**Library:**

`libc`

**Description:**

The `setrlimit()` function sets the limits on the consumption of a variety of system resources by a process and each process it creates. The `setrlimit64()` function is a 64-bit version of `setrlimit()`.

Each call to `setrlimit()` identifies a specific resource to be operated upon as well as a resource limit. A resource limit is a pair of values:

one specifying the current (soft) limit, the other a maximum (hard) limit.

A process can change soft limits to any value that's less than or equal to the hard limit. A process may (irreversibly) lower its hard limit to any value that's greater than or equal to the soft limit. Only a process with an effective user ID of superuser can raise a hard limit. Both hard and soft limits can be changed in a single call to *setrlimit()*, subject to the constraints described above. Limits may have an "infinite" value of RLIM\_INFINITY.



RLIM\_INFINITY is a special value, and it's actual numerical value doesn't represent a valid VM/AS size in bytes.

The possible resources, their descriptions, and the actions taken when the current limit is exceeded are summarized below:

| Resource    | Description                                                                                                                     | Action                                                                                                                          |
|-------------|---------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| RLIMIT_CORE | The maximum size, in bytes, of a core file that may be created by a process. A limit of 0 prevents the creation of a core file. | The writing of a core file terminates at this size.                                                                             |
| RLIMIT_CPU  | The maximum amount of CPU time, in seconds, used by a process. This is a soft limit only.                                       | SIGXCPU is sent to the process. If the process is holding or ignoring SIGXCPU, the behavior is defined by the scheduling class. |
| RLIMIT_DATA | The maximum size of a process's heap in bytes                                                                                   | The <i>brk()</i> function fails with <i>errno</i> set to ENOMEM.                                                                |

*continued...*

| <b>Resource</b> | <b>Description</b>                                                                                                                                                        | <b>Action</b>                                                                                                                                                                                     |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RLIMIT_FSIZE    | The maximum size of a file in bytes that may be created by a process. A limit of 0 prevents the creation of a file.                                                       | The SIGXFSZ signal is sent to the process. If the process is holding or ignoring SIGXFSZ, continued attempts to increase the size of a file beyond the limit fail with <i>errno</i> set to EFBIG. |
| RLIMIT_NOFILE   | One more than the maximum value that the system may assign to a newly created descriptor. This limit constrains the number of file descriptors that a process may create. |                                                                                                                                                                                                   |

*continued...*

| Resource     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Action                                                                                                                                                                                                                                         |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RLIMIT_STACK | <p>The maximum size of a process's stack in bytes. The system will not automatically grow the stack beyond this limit.</p> <p>Within a process, <i>setrlimit()</i> increases the limit on the size of your stack, but doesn't move current memory segments to allow for that growth. To guarantee that the process stack can grow to the limit, the limit must be altered prior to the execution of the process in which the new stack size is to be used.</p> <p>Within a multithreaded process, <i>setrlimit()</i> has no impact on the stack size limit for the calling thread if the calling thread isn't the main thread. A call to <i>setrlimit()</i> for RLIMIT_STACK impacts only the main thread's stack, and should be made only from the main thread, if at all.</p> | <p>The SIGSEGV signal is sent to the process. If the process is holding or ignoring SIGSEGV, or is catching SIGSEGV and hasn't made arrangements to use an alternate stack, the disposition of SIGSEGV is set to SIG_DFL before it's sent.</p> |
| RLIMIT_VMEM  | <p>The maximum size of a process's mapped address space in bytes.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>If the limit is exceeded, the <i>brk()</i>, <i>mmap()</i> and <i>sbrk()</i> functions fail with <i>errno</i> set to ENOMEM. In addition, the automatic stack growth fails with the effects outlined above.</p>                              |

*continued...*



| Resource  | Description          | Action               |
|-----------|----------------------|----------------------|
| RLIMIT_AS | Same as RLIMIT_VMEM. | Same as RLIMIT_VMEM. |

Because limit information is stored in the per-process information, the shell builtin `ulimit` command must directly execute this system call if it's to affect all future processes created by the shell.

The values of the current limit of the following resources affect these parameters:

| Resource      | Parameter |
|---------------|-----------|
| RLIMIT_FSIZE  | FCHR_MAX  |
| RLIMIT_NOFILE | OPEN_MAX  |

When using the *setrlimit()* function, if the requested new limit is `RLIM_INFINITY`, there's no new limit; otherwise, if the requested new limit is `RLIM_SAVED_MAX`, the new limit is the corresponding saved hard limit; otherwise, if the requested new limit is `RLIM_SAVED_CUR`, the new limit is the corresponding saved soft limit; otherwise, the new limit is the requested value. In addition, if the corresponding saved limit can be represented correctly in an object of type `rlim_t`, then it's overwritten with the new limit.

The result of setting a limit to `RLIM_SAVED_MAX` or `RLIM_SAVED_CUR` is unspecified unless a previous call to *getrlimit()* returned that value as the soft or hard limit for the corresponding resource limit.

A limit whose value is greater than `RLIM_INFINITY` is permitted.

The *exec\** family of functions also cause resource limits to be saved.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EFAULT The *rlp* argument points to an illegal address.
- EINVAL An invalid resource was specified, the new *rlim\_cur* exceeds the new *rlim\_max*, or the limit specified can't be lowered because current usage is already higher than the limit.
- EPERM The limit specified to *setrlimit()* would've raised the maximum limit value, and the effective user of the calling process isn't the superuser.

## Classification:

*setrlimit()* is standard Unix; *setrlimit64()* is for large-file support

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*brk()*, *execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *fork()*, *getdtablesize()*, *getrlimit()*, *getrlimit64()*, *malloc()*, *open()*, *signal()*, *sysconf()*

**Synopsis:**

```
#include <netdb.h>

void setservent(int stayopen);
```

**Arguments:**

*stayopen*      Nonzero if you don't want the database to be closed after each call to *getservbyname()* or *getservbyport()*.

**Library:**

**libsocket**

**Description:**

The *setservent()* function opens and rewinds the network services database file. If the *stayopen* flag is nonzero, the network services database won't be closed after each call to *getservbyname()* or *getservbyport()*.

**Files:**

**/etc/services**  
Network services database file.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

**Caveats:**

This function uses static data; if you need the data for future use, you should copy it before any subsequent calls overwrite it.

**See also:**

*endservent()*, *getservbyname()*, *getservbyport()*, *getservent()*,  
**servent**

*/etc/services* in the *Utilities Reference*.

## Synopsis:

```
#include <unistd.h>

pid_t setsid(void);
```

## Library:

libc

## Description:

The *setsid()* function creates a new session with the calling process becoming the process group leader with no controlling terminal. The process group ID of the calling process is set to the process ID of the calling process. The calling process is the only process in the new process group, and is also the only process in the new session.

If the calling process is already a process group leader, a new session isn't created and an error is returned.

## Returns:

The new process group ID for the calling process, or -1 if an error occurred (*errno* is set).

## Errors:

EPERM      The calling process is already a process group leader, or the process group ID of a process other than the calling process matches the process ID of the calling process.

## Examples:

```
/*
 * You can only become a session leader if you are not
 * a process group leader that is the default for a
 * command run from the shell.
 */

#include <stdio.h>
#include <sys/types.h>
```

```
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 if(fork())
 {
 if(setsid() == -1)
 perror("parent: setsid");
 else
 printf("parent: I am a session leader\n");
 }
 else
 {
 if(setsid() == -1)
 perror("child: setsid");
 else
 printf("child: I am a session leader\n");
 }
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *getsid()*, *setpgid()*

**Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

int setsockopt(int s,
 int level,
 int optname,
 const void * optval,
 socklen_t optlen);
```

**Arguments:**

- |                |                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>s</i>       | The file descriptor of the socket that the option is to be applied on, as returned by <i>socket()</i> .                                                                                                                                                                                                                                                                                          |
| <i>level</i>   | The protocol layer that the option is to be applied to. In most cases, it's a socket-level option and is indicated by SOL_SOCKET.                                                                                                                                                                                                                                                                |
| <i>optname</i> | The option for the socket file descriptor.                                                                                                                                                                                                                                                                                                                                                       |
| <i>optval</i>  | A pointer to the value of the option (in most cases, whether the option is to be turned on or off). If no option value is to be returned, <i>optval</i> may be NULL.<br><br>Most socket-level options use an <b>int</b> parameter for <i>optval</i> . Others, such as the SO_LINGER, SO_SNDTIMEO, and SO_RCVTIMEO options, use structures that also let you get data associated with the option. |
| <i>optlen</i>  | A pointer to the length of the value of the option. This argument is a value-result parameter; you should initialize it to indicate the size of the buffer pointed to by <i>optval</i> .                                                                                                                                                                                                         |

## Library:

libsocket

## Description:

The *setsockopt()* function sets the options associated with a socket.

See *getsockopt()* for more detailed information.

## Returns:

0 Success.

-1 An error occurred (*errno* is set).

## Errors:

EBADF Invalid file descriptor *s*.

EDOM Value was set out of range.

EFAULT The address pointed to by *optval* isn't in a valid part of the process address space.

EINVAL No *optval* value was specified.

ENOPROTOOPT

The option is unknown at the level indicated.

## Classification:

Standard Unix, POSIX 1003.1-2001

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



**See also:**

ICMP, IP, TCP, and UDP protocols

*getsockopt()*, *socket()*

## ***setspent()***

© 2004, QNX Software Systems Ltd.

*Rewind the shadow password database file*

### **Synopsis:**

```
#include <sys/types.h>
#include <shadow.h>

void setspent(void);
```

### **Library:**

libc

### **Description:**

The *setspent()* function rewinds to the start of the shadow password database file. It's provided for programs that make multiple lookups in the database (using the *getspnam()* call) to avoid opening and closing the shadow password database for each access.

### **Classification:**

Unix

#### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

### **See also:**

*fgetspent() endspent(), getspnam(), getspent() putspent()*

## Synopsis:

```
#include <stdlib.h>

char *setstate(const char *state);
```

## Arguments:

*state*     A pointer to the state array that you want to use.

## Library:

`libc`

## Description:

Once the state of the pseudo-random number generator has been initialized, *setstate()* allows switching between state arrays. The array defined by the *state* argument is used for further random-number generation until *initstate()* is called or *setstate()* is called again. The *setstate()* function returns a pointer to the previous state array.

This function is used in conjunction with the following:

|                    |                                                             |
|--------------------|-------------------------------------------------------------|
| <i>initstate()</i> | Initialize the state of the pseudo-random number generator. |
| <i>random()</i>    | Generate a pseudo-random number using a default state.      |
| <i>srandom()</i>   | Set the seed used by the pseudo-random number generator.    |

After initialization, you can restart a state array at a different point in one of two ways:

- Call *initstate()* with the desired seed, state array, and size of the array.

- Call *setstate()* with the desired state, then call *srandom()* with the desired seed. The advantage of using both of these functions is that the size of the state array doesn't have to be saved once it's initialized.

**Returns:**

A pointer to the previous state array, or NULL if an error occurred.

**Examples:**

See *initstate()*.

**Classification:**

Standard Unix

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*drand48()*, *initstate()*, *rand()*, *random()*, *srand()*, *srandom()*

### **Synopsis:**

```
#include <sys/time.h>

int settimeofday(const struct timeval *when,
 void *not_used);
```

### **Arguments:**

*when*            A pointer to a `timeval` structure that specifies the time that you want to set. The `struct timeval` contains the following members:

- `long tv_sec` — the number of seconds since the start of the Unix Epoch.
- `long tv_usec` — the number of microseconds.

*not\_used*        This pointer must be NULL or the behavior of `settimeofday()` is unspecified. This argument is provided only for backwards compatibility.

### **Library:**

`libc`

### **Description:**

This function sets the time and date to the values stored in the structure pointed to by *when*.

### **Returns:**

0, or -1 if an error occurred (*errno* is set).

### **Errors:**

EFAULT        An error occurred while accessing the *when* buffer.

EPERM         The calling process doesn't have superuser permissions.

## Classification:

Legacy Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The *settimeofday()* function is provided for compatibility with existing Unix code. You shouldn't use it in new code — use *clock\_settime()* instead.

## See also:

*asctime()*, *asctime\_r()*, *clock\_gettime()*, *clock\_settime()*, *ctime()*, *ctime\_r()*, *difftime()*, *gettimeofday()*, *gmtime()*, *gmtime\_r()*, *localtime()*, *localtime\_r()*, *time()*

## Synopsis:

```
#include <unistd.h>

int setuid(uid_t uid);
```

## Arguments:

*uid*     The user ID that you want to use for the process.

## Library:

`libc`

## Description:

The *setuid()* function lets the calling process set the real, effective and saved user IDs based on the following:

- If the process is the superuser, *setuid()* sets the real user ID, effective user ID and saved user ID to *uid*.
- If the process isn't the superuser, but *uid* is equal to the real user ID or saved set-user ID, *setuid()* sets the effective user ID to *uid*; the real and saved user IDs aren't changed.



---

If a set-UID process sets its effective user ID to its real user ID, it can still set its effective user ID back to the saved set-UID.

---

If you wish to change only the effective user ID, and even if you are the superuser, you should consider using the *seteuid()* function.

The “superuser” is defined as any process with an effective user ID of 0, or an effective user ID of `root`.

**Returns:**

0 for success, or -1 if an error occurs (*errno* is set).

**Errors:**

EINVAL     The value of *uid* is out of range.

EPERM     The process isn't the superuser, and *uid* doesn't match the real user ID or saved set-user ID.

**Examples:**

```
/*
 * This process sets its userid to 0 (root)
 */
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 uid_t ould;

 ould = getuid();
 if(setuid(0) == -1) {
 perror("setuid");
 return EXIT_FAILURE;
 }

 printf("userid %d switched to 0\n", ould);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

Cancellation point   No

*continued...*



**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*errno, getuid(), setegid(), seteuid(), setgid()*

## ***setutent()***

© 2004, QNX Software Systems Ltd.

*Return to the beginning of the user-information file*

### **Synopsis:**

```
#include <utmp.h>

void setutent(void);
```

### **Library:**

libc

### **Description:**

The *setutent()* function resets the input stream to the beginning of the file specified in *\_PATH\_UTMP*. Do this before each search for a new entry if you want the entire file to be examined.

### **Files:**

*\_PATH\_UTMP*  
Specifies the user information file.

### **Classification:**

Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*endutent()*, *getutent()*, *getutid()*, *getutline()*, *pututline()*, **utmp**,  
*utmpname()*

**login** in the *Utilities Reference*

## ***setvbuf()***

© 2004, QNX Software Systems Ltd.

*Associate a buffer with a stream*

### **Synopsis:**

```
#include <stdio.h>

int setvbuf(FILE *fp,
 char *buf,
 int mode,
 size_t size);
```

### **Arguments:**

*fp*            The stream that you want to associate with a buffer.

*buffer*        NULL, or a pointer to the buffer; see below.

*mode*          How you want the stream to be buffered:

- `.IOFBF` — input and output are fully buffered.
- `.IOLBF` — output is line buffered (i.e. the buffer is flushed when a newline character is written, when the buffer is full, or when input is requested).
- `.IONBF` — input and output are completely unbuffered.

*size*          The size of the buffer.

### **Library:**

`libc`

### **Description:**

The *setvbuf()* function associates a buffer with the stream designated by *fp*. If you want to call *setvbuf()*, you must call it after opening the stream, but before doing any reading, writing, or seeking.

If *buf* isn't NULL, the buffer it points to is used instead of an automatically allocated buffer.

**Returns:**

- |        |                                                                                                                     |
|--------|---------------------------------------------------------------------------------------------------------------------|
| 0      | Success.                                                                                                            |
| EINVAL | The <i>mode</i> argument isn't valid.                                                                               |
| ENOMEM | The <i>buf</i> argument is NULL, <i>size</i> isn't 0, and there isn't enough memory available to allocate a buffer. |

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char *buf;
 FILE *fp;

 fp = fopen("file", "r");
 buf = malloc(1024);
 setvbuf(fp, buf, _IOFBF, 1024);

 /* work with fp */
 ...

 fclose(fp);

 /* This is OUR buffer, so we have
 * to free it. Do that AFTER
 * you've closed the file.
 */

 free(buf);
 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fopen(), setbuf()*

## Synopsis:

```
#include <malloc.h>

void _sfree(void *ptr,
 size_t size);
```

## Arguments:

*ptr*      NULL, or a pointer to the block of memory that you want to free.

*size*     The number of bytes to deallocate.

## Library:

`libc`

## Description:

When the value of the argument *ptr* is NULL, the *\_sfree()* function does nothing; otherwise, the *\_sfree()* function deallocates *size* bytes of memory located by the argument *ptr*, which was previously returned by a call to the appropriate version of *\_scalloc()* or *\_smalloc()*. After the call, the freed block is available for allocation.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Calling `_sfree()` on a pointer already deallocated by a call to `_sfree()` could corrupt the memory allocator's data structures.

The *size must* match the size of the allocated block.

**See also:**

`calloc()`, `free()`, `realloc()`, `_scalloc()`, `_smalloc()`, `_srealloc()`



**Synopsis:**

```
#include <sys/mman.h>

int shm_ctl(int fd,
 int flags,
 _uint64 paddr,
 _uint64 size);
```

**Arguments:**

- fd* The file descriptor that's associated with the shared memory object, as returned by *shm\_open()*.
- flags* One or more of the following bits, defined in *<sys/mman.h>*:
- SHMCTL\_ANON — grow the object to be *size* bytes.
  - SHMCTL\_PHYS — assign the physical address, *paddr*, to the object.
  - SHMCTL\_GLOBAL — a hint that any mapping to the object could be global across all processes.
  - SHMCTL\_PRIV — a hint that a mapping of this object may require privileged access.
  - SHMCTL\_LOWERPROT — a hint that the system may map this object in such a way that it trades lower memory protection for better performance.
- paddr* A physical address to assign to the object, if you set SHMCTL\_PHYS in *flags*.
- size* The new size of the object, in bytes, if you set SHMCTL\_ANON in *flags*.

## Library:

libc

## Description:

The *shm\_ctl()* function modifies the attributes of the shared memory object identified by the handle, *fd*. This handle is the value returned by *shm\_open()*.



---

The combination SHMCTL\_ANON | SHMCTL\_PHYS has the same behavior as for *mmap()*: it indicates that you want physically contiguous RAM to be allocated for the object.

---

## Returns:

0      Success.

-1     An error occurred (*errno* is set).

## Errors:

EINVAL     An invalid combination of flags was specified, or the shared memory object is already “special.”

## Examples:

The following examples go together. Run **sharephyscreator**, followed by **sharephysuser**.

The **sharephyscreator** process maps in an area of physical memory and then overlays it with a shared memory object. The **sharephysuser** process opens that shared memory object in order to access the physical memory.

```
/*
 * sharephyscreator.c
 *
 * This maps in an area of physical memory and then
 * overlays it with a shared memory object. This way, another process
 * can open that shared memory object in order to access the physical
 * memory. The other process in this case is sharephysuser.
```

```

*
* Note that the size and address that you pass to shm_ctl() must be
* even multiples of the page size (sysconf(_SC_PAGE_SIZE)).
*
* For VGA color text mode video memory:
* sharephyscreator /wally b8000
* Note that for VGA color text mode video memory, each character
* is followed by an attribute byte. Here we just use a space.
*/

#include <errno.h>
#include <fcntl.h>
#include <inttypes.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h>
#include <sys/mman.h>
#include <sys/neutrino.h>
#include <sys/stat.h>

char *programe = "sharephyscreator";

main(int argc, char *argv[])
{
 char *text = "H e l l o w o r l d ! ";
 int fd, memsize;
 char *ptr, *name;
 uint64_t physaddr;

 if (argc != 3) {
 printf("use: sharephyscreator shared_memory_object_name physical_address_in_hex\n");
 printf("Example: sharephyscreator wally b8000\n");
 exit(EXIT_FAILURE);
 }
 name = argv[1];
 physaddr = atoh(argv[2]);
 memsize = sysconf(_SC_PAGE_SIZE); /* this should be enough
 for our string */

 /* map in the physical memory */

 ptr = mmap_device_memory(0, memsize, PROT_READ|PROT_WRITE, 0, physaddr);
 if (ptr == MAP_FAILED) {
 printf("%s: mmap_device_memory for physical address %llx failed: %s\n",
 programe, physaddr, strerror(errno));
 exit(EXIT_FAILURE);
 }

 /* open the shared memory object, create it if it doesn't exist */

 fd = shm_open(name, O_RDWR | O_CREAT, 0);
 if (fd == -1) {
 printf("%s: error creating the shared memory object '%s': %s\n",
 programe, name, strerror(errno));
 exit(EXIT_FAILURE);
 }

 /* overlay the shared memory object onto the physical memory */

```

```

if (shm_ctl(fd, SHMCTL_PHYS, physaddr, memsize) == -1) {
 printf("%s: shm_ctl failed: %s\n", progname, strerror(errno));
 close(fd);
 munmap(ptr, memsize);
 shm_unlink(name);
 exit(EXIT_FAILURE);
}
strcpy(ptr, text); /* write to the shared memory */

printf("\n%s: Physical memory mapped in, shared memory overlaid onto it.\n"
 "%s: Wrote '%s' to physical memory.\n"
 "%s: Sleeping for 20 seconds. While this program is sleeping\n"
 "%s: run 'sharephysuser %s %d'.\n",
 progname, progname, ptr, progname, progname, name,
 strlen(text)+1);
sleep(20);

printf("%s: Woke up. Cleaning up and exiting ...\n", progname);

close(fd);
munmap(ptr, memsize);
shm_unlink(name);
}

```

The following is meant to be run with `sharephyscreator`.

```

/*
 * sharephysuser.c
 *
 * This one is meant to be run in tandem with sharephyscreator.
 *
 * Run it as: sharephysuser shared_memory_object_name length
 * Example: sharephysuser wally 49
 *
 */

#include <errno.h>
#include <fcntl.h>
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <unistd.h>
#include <sys/mman.h>
#include <sys/neutrino.h>
#include <sys/stat.h>

char *progname = "sharephysuser";

main(int argc, char *argv[])
{
 int fd, len, i;
 char *ptr, *name;

```

```

if (argc != 3) {
 fprintf(stderr, "use: sharephysuser shared_memory_object_name length\n");
 fprintf(stderr, "Example: sharephysuser wally 49\n");
 exit(EXIT_FAILURE);
}
name = argv[1];
len = atoi(argv[2]);

/* open the shared memory object */

fd = shm_open(name, O_RDWR, 0);
if (fd == -1) {
 fprintf(stderr, "%s: error opening the shared memory object '%s': %s\n",
 progname, name, strerror(errno));
 exit(EXIT_FAILURE);
}

/* get a pointer to a piece of the shared memory, note that we
 only map in the amount we need to */

ptr = mmap(0, len, PROT_READ | PROT_WRITE, MAP_SHARED, fd, 0);
if (ptr == MAP_FAILED) {
 fprintf(stderr, "%s: mmap failed: %s\n", progname, strerror(errno));
 exit(EXIT_FAILURE);
}

printf("%s: reading the text: ", progname);
for (i = 0; i < len; i++)
 printf("%c", ptr[i]);
printf("\n");

close(fd);
munmap(ptr, len);
}

```

**Classification:**

QNX Neutrino

**Safety**

Cancellation point Yes

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*mmap()*, *munmap()*, *mprotect()*, *shm\_open()*, *shm\_unlink()*

**Synopsis:**

```
#include <fcntl.h>
#include <sys/mman.h>

int shm_open(const char * name,
 int oflag,
 mode_t mode);
```

**Arguments:**

- name*     The name of the shared memory object that you want to open; see below.
- oflag*     A combination of the following bits (defined in <fcntl.h>):
- O\_RDONLY — open for read access only.
  - O\_RDWR — open for read and write access.
  - O\_CREAT — if the shared memory object exists, this flag has no effect, except as noted under O\_EXCL below. Otherwise, the shared memory object is created, and its permissions are set in accordance with the value of *mode* and the file mode creation mask of the process.
  - O\_EXCL — if O\_EXCL and O\_CREAT are set, then *shm\_open()* fails if the shared memory segment exists. The check for the existence of the shared memory object, and the creation of the object if it doesn't exist, are atomic with respect to other processes executing *shm\_open()*, naming the same shared memory object with O\_EXCL and O\_CREAT set.
  - O\_TRUNC — if the shared memory object exists, and it's successfully opened O\_RDWR, the object is truncated to zero length and the mode and owner are unchanged.
- mode*     The permission bits for the memory object are set to the value of *mode*, except those bits set in the process's file creation mask. For more information, see *umask()*, and "Access permissions" in the documentation for *stat()*.

## Library:

libc

## Description:

The *shm\_open()* function returns a file descriptor that's associated with the shared "memory object" specified by *name*. This file descriptor is used by other functions to refer to the shared memory object (for example, *mmap()*, *mprotect()*). The `FD_CLOEXEC` file descriptor flag in *fcntl()* is set for this file descriptor.

The *name* argument is interpreted as follows:

| <i>name</i>            | Pathname space entry      |
|------------------------|---------------------------|
| <i>entry</i>           | <i>CWD/entry</i>          |
| <i>/entry</i>          | <i>/dev/shmem/entry</i>   |
| <i>entry/newentry</i>  | <i>CWD/entry/newentry</i> |
| <i>/entry/newentry</i> | <i>/entry/newentry</i>    |

where *CWD* is the current working directory for the program at the point that it calls *mq\_open()*.



If you want to open a shared memory object on another node, you have to specify the name as */net/node/shmem\_location*.

The state of the shared memory object, including all data associated with it, persists until the shared memory object is unlinked and all other references are gone.

## Returns:

A nonnegative integer, which is the lowest numbered unused file descriptor, or -1 if an error occurred (*errno* is set).



**Errors:**

|              |                                                                                                                                                                                                                   |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Permission to create the shared memory object is denied.<br><br>The shared memory object exists and the permissions specified by <i>oflag</i> are denied, or O_TRUNC is specified and write permission is denied. |
| EEXIST       | O_CREAT and O_EXCL are set, and the named shared memory object already exists.                                                                                                                                    |
| EINTR        | The <i>shm_open()</i> call was interrupted by a signal.                                                                                                                                                           |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                                                    |
| EMFILE       | Too many file descriptors are currently in use by this process.                                                                                                                                                   |
| ENAMETOOLONG | The length of the <i>name</i> argument exceeds NAME_MAX.                                                                                                                                                          |
| ENFILE       | Too many shared memory objects are currently open in the system.                                                                                                                                                  |
| ENOENT       | O_CREAT isn't set, and the named shared memory object doesn't exist, or O_CREAT is set and either the name prefix doesn't exist or the <i>name</i> argument points to an empty string.                            |
| ENOSPC       | There isn't enough space to create the new shared memory object.                                                                                                                                                  |
| ENOSYS       | The <i>shm_open()</i> function isn't supported by this implementation.                                                                                                                                            |

**Examples:**

This example sets up a shared memory object, but doesn't really do anything with it:

```
#include <stdio.h>
#include <string.h>
#include <fcntl.h>
#include <errno.h>
#include <stdlib.h>
#include <unistd.h>
#include <limits.h>
#include <sys/mman.h>

int main(int argc, char** argv)
{
 int fd;
 unsigned* addr;

 /*
 * In case the unlink code isn't executed at the end
 */
 if(argc != 1) {
 shm_unlink("/bolts");
 return EXIT_SUCCESS;
 }

 /* Create a new memory object */
 fd = shm_open("/bolts", O_RDWR | O_CREAT, 0777);
 if(fd == -1) {
 fprintf(stderr, "Open failed:%s\n",
 strerror(errno));
 return EXIT_FAILURE;
 }

 /* Set the memory object's size */
 if(ftruncate(fd, sizeof(*addr)) == -1) {
 fprintf(stderr, "ftruncate: %s\n",
 strerror(errno));
 return EXIT_FAILURE;
 }

 /* Map the memory object */
 addr = mmap(0, sizeof(*addr),
 PROT_READ | PROT_WRITE,
 MAP_SHARED, fd, 0);
 if(addr == MAP_FAILED) {
 fprintf(stderr, "mmap failed: %s\n",
 strerror(errno));
 return EXIT_FAILURE;
 }

 printf("Map addr is 0x%08x\n", addr);

 /* Write to shared memory */
```

```
 *addr = 1;

 /*
 * The memory object remains in
 * the system after the close
 */
 close(fd);

 /*
 * To remove a memory object
 * you must unlink it like a file.
 *
 * This may be done by another process.
 */
 shm_unlink("/bolts");

 return EXIT_SUCCESS;
}
```

This example uses a shared memory object to share data with a forked process:

```
#include <stdio.h>
#include <string.h>
#include <fcntl.h>
#include <errno.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/mman.h>

main(int argc, char * argv[])
{
 int fd;
 unsigned *addr;

 /*
 * In case the unlink code isn't executed at the end
 */
 if (argc != 1) {
 shm_unlink("/bolts");
 return EXIT_SUCCESS;
 }

 /* Create a new memory object */
 fd = shm_open("/bolts", O_RDWR | O_CREAT, 0777);
 if (fd == -1) {
 fprintf(stderr, "Open failed : %s\n",
 strerror(errno));
 return EXIT_FAILURE;
 }
}
```

```
 }

 /* Set the memory object's size */
 if (ftruncate(fd, sizeof(*addr)) == -1) {
 fprintf(stderr, "ftruncate : %s\n", strerror(errno));
 return EXIT_FAILURE;
 }

 /* Map the memory object */
 addr = mmap(0, sizeof(*addr), PROT_READ | PROT_WRITE,
 MAP_SHARED, fd, 0);
 if (addr == MAP_FAILED) {
 fprintf(stderr, "mmap failed:%s\n", strerror(errno));
 return EXIT_FAILURE;
 }

 printf("Map addr is %6.6X\n", addr);
 printf("Press break to stop.\n");
 sleep(3); /* So you can read above message */

 /*
 * We unlink so object goes away on last close.
 */
 shm_unlink("/bolts");

 *addr = '0';
 if (fork())
 for (;;)
 if (*addr == '0')
 putc(*addr = '1', stderr);
 else
 sched_yield();
 else
 for (;;)
 if (*addr == '1')
 putc(*addr = '0', stderr);
 else
 sched_yield();
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*fcntl()*, *ftruncate()*, *mmap()*, *munmap()*, *mprotect()*, *open()*, *shm\_ctl()*,  
*shm\_unlink()*, *sysconf()*

## ***shm\_unlink()***

© 2004, QNX Software Systems Ltd.

*Remove a shared memory object*

### **Synopsis:**

```
#include <sys/mman.h>

int shm_unlink(const char * name);
```

### **Arguments:**

*name*     The name of the shared memory object that you want to remove.

### **Library:**

`libc`

### **Description:**

The *shm\_unlink()* function removes the name of the shared memory object specified by *name*. After removing the name, you can't use *shm\_open()* to access the object.

This function doesn't affect any references to the shared memory object (i.e. file descriptors or memory mappings). If more than one reference to the shared memory object exists, then the link count is decremented, but the shared memory segment isn't actually removed until you remove all open and map references to it.

### **Returns:**

0        Success.

-1       An error occurred (*errno* is set).

### **Errors:**

EACCES    Permission to unlink the shared memory object is denied.

ELOOP    Too many levels of symbolic links or prefixes.

ENAMETOOLONG

The length of the *name* argument exceeds NAME\_MAX.

ENOENT

The named shared memory object doesn't exist, or the *name* argument points to an empty string.

ENOSYS

The *shm\_unlink()* function isn't supported by this implementation.

**Examples:**

See *shm\_open()*.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*mmap()*, *munmap()*, *mprotect()*, *shm\_ctl()*, *shm\_open()*

## ***shutdown()***

© 2004, QNX Software Systems Ltd.

*Shut down part of a full-duplex connection*

---

### **Synopsis:**

```
#include <sys/socket.h>

int shutdown(int s,
 int how);
```

### **Arguments:**

*s*        A descriptor for the socket, as returned by *socket()*.  
*how*     How you want to shut down the connection:

| <b>If <i>how</i> is:</b> | <b>The TCP/IP manager won't allow:</b> |
|--------------------------|----------------------------------------|
|--------------------------|----------------------------------------|

---

|         |                  |
|---------|------------------|
| SHUT_RD | Further receives |
|---------|------------------|

|         |               |
|---------|---------------|
| SHUT_WR | Further sends |
|---------|---------------|

|           |                            |
|-----------|----------------------------|
| SHUT_RDWR | Further sends and receives |
|-----------|----------------------------|

### **Library:**

`libsocket`

### **Description:**

The *shutdown()* call shuts down all or part of a full-duplex connection on the socket associated with *s*.

### **Returns:**

0        Success.  
-1      An error occurred (*errno* is set).



**Errors:**

EBADF      Invalid descriptor *s*.

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*close()*, *connect()*, *socket()*

## ***sigaction()***

© 2004, QNX Software Systems Ltd.

*Examine or specify the action associated with a signal*

### **Synopsis:**

```
#include <signal.h>

int sigaction(int sig,
 const struct sigaction * act,
 struct sigaction * oact);
```

### **Arguments:**

- sig*      The signal number (defined in <**signal.h**>). For more information, see “POSIX signals” in the documentation for *SignalAction()*.
- act*      NULL, or a pointer to a **sigaction** structure that specifies how you want to modify the action for the given signal. For more information about this structure, see below.
- oact*     NULL, or a pointer to a **sigaction** structure that the function fills with information about the current action for the signal.

### **Library:**

**libc**

### **Description:**

You can use *sigaction()* to examine or specify (or both) the action that's associated with a specific signal:

- If *act* isn't NULL, the specified signal is modified.
- If *oact* isn't NULL, the previous action is stored in the structure it points to.

The structure **sigaction** contains the following members:

```
void (*sa_handler)();
```

Address of a signal handler or action for nonqueued signals.

**void** (\**sa\_sigaction*)(**int** *signo*, **siginfo\_t**\* *info*, **void**\* *other*);  
Address of a signal handler or action for queued signals.

**sigset\_t** *sa\_mask*

An additional set of signals to be masked (blocked) during execution of the signal-catching function.

**int** *sa\_flags*

Special flags to affect behavior of the signal:

- SA\_NOCLDSTOP is only used when the signal is SIGCHLD. It tells the system not to generate a SIGCHLD on the parent for children who stop via SIGSTOP.
- SA\_SIGINFO tells the OS to queue this signal. The default is not to queue a signal delivered to a process. If a signal isn't queued, setting the same signal multiple times on a process or thread before it runs results in only the last signal's being delivered. If you set SA\_SIGINFO, the signals are queued and they're all delivered.

The *sa\_handler* and *sa\_sigaction* members of *act* are implemented as a **union** and share common storage. They differ only in their prototypes, with *sa\_handler* being used for POSIX 1003.1a signals and *sa\_sigaction* being used for POSIX 1003.1b queued realtime signals. The values stored using either name can be one of:

*function*      The address of a signal catching function. See below for details.

SIG\_DFL      This sets the signal to the default action:

- SIGCHLD, SIGIO, SIGURG and SIGWINCH — ignore the signal (SIG\_IGN).
- SIGSTOP — stop the process.
- SIGCONT — continue the program.

- All other signals — kill the process.

**SIG\_IGN** This ignores the signal. Setting **SIG\_IGN** for a signal that's pending discards all pending signals, whether it's blocked or not. New signals are discarded. If you ignore **SIGCHLD**, your process's children don't enter the zombie state and you're unable to wait on their death using *wait()* or *waitpid()*.

The *function* member of *sa\_handler* or *sa\_sigaction* is always invoked with the following arguments:

```
void handler(int signo, siginfo_t *info, void *other)
```

If you have an old-style signal handler of the form:

```
void handler(int signo)
```

the extra arguments are still placed by the kernel, but the function simply ignores them.

While in the handler, *signo* is masked, preventing nested signals of the same type. In addition, any signals set in the *sa\_mask* member of *act* are also ORed into the mask. When the handler returns through a normal return, the previous mask is restored and any pending and now unmasked signals are acted on. You return to the point in the program where it was interrupted. If the thread was blocked in the kernel when the interruption occurred, the kernel call returns with an **EINTR** (see *ChannelCreate()* and *SyncMutexLock()* for exceptions to this).

The **siginfo\_t** structure of the *function* in *sa\_handler* or *sa\_sigaction* contains at least the following members:

**int** *si\_signo* The signal number, which should match the *signo* argument to the handler.

**int** *si\_code* A signal code, provided by the generator of the signal:

- `SI_USER` — the *kill()* function generated the signal.
- `SI_QUEUE` — the *sigqueue()* function generated the signal.
- `SI_TIMER` — a timer generated the signal.
- `SI_ASYNCIO` — asynchronous I/O generated the signal.
- `SI_MSGQ` — POSIX (not QNX) messages queues generated the signal.

**union signal** *si\_value*

A value associated with the signal, provided by the generator of the signal.

You can't ignore or catch `SIGKILL` or `SIGSTOP`.

Signal handlers and actions are defined for the process and affect all threads in the process. For example, if one thread ignores a signal, then all threads ignore the signal.

You can target a signal at a thread, process, or process group (see *SignalKill()*). When targeted at a process, at most one thread receives the signal. This thread must have the signal unblocked (see *SignalProcmask()*) to be a candidate for receiving it. All synchronously generated signals (e.g. `SIGSEGV`) are always delivered to the thread that caused them.



---

If you use *longjmp()* to return from a signal handler, the signal remains masked. You can use *siglongjmp()* to restore the mask to the state saved by a previous call to *sigsetjmp()*.

---

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

|        |                                                                            |
|--------|----------------------------------------------------------------------------|
| EAGAIN | Insufficient system resources are available to set up the signal's action. |
| EFAULT | A fault occurred trying to access the buffers provided.                    |
| EINVAL | The signal <i>signo</i> isn't valid.                                       |

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>

int main(void)
{
 extern void handler();
 struct sigaction act;
 sigset_t set;

 sigemptyset(&set);
 sigaddset(&set, SIGUSR1);
 sigaddset(&set, SIGUSR2);

 /*
 * Define a handler for SIGUSR1 such that when
 * entered both SIGUSR1 and SIGUSR2 are masked.
 */
 act.sa_flags = 0;
 act.sa_mask = set;
 act.sa_handler = &handler;
 sigaction(SIGUSR1, &act, NULL);

 kill(getpid(), SIGUSR1);

 /* Program will terminate with a SIGUSR2 */
 return EXIT_SUCCESS;
}

void handler(signo)
{
 static int first = 1;

 if(first) {
 first = 0;
 kill(getpid(), SIGUSR1); /* Prove signal masked */
 kill(getpid(), SIGUSR2); /* Prove signal masked */
 }
}
```

```
 }
}

/*
 * - SIGUSR1 is set from main(), handler() is called.
 * - SIGUSR1 and SIGUSR2 are set from handler().
 * - however, signals are masked until we return to main().
 * - returning to main() unmask SIGUSR1 and SIGUSR2.
 * - pending SIGUSR1 now occurs, handler() is called.
 * - pending SIGUSR2 now occurs. Since we don't have
 * a handler for SIGUSR2, we are killed.
 */
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *kill()*, *raise()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *signal()*, *SignalAction()*, *SignalKill()*, *sigpending()*, *sigprocmask()*

## ***sigaddset()***

© 2004, QNX Software Systems Ltd.

*Add a signal to a set*

### **Synopsis:**

```
#include <signal.h>

int sigaddset(sigset_t *set,
 int signo);
```

### **Arguments:**

*set*        A pointer to the `sigset_t` object that you want to add the signal to.

*signo*     The signal that you want to add. For more information, see “POSIX signals” in the documentation for *SignalAction()*.

### **Library:**

`libc`

### **Description:**

The *sigaddset()* function adds *signo* to the set pointed to by *set*.

### **Returns:**

0        Success.

-1      An error occurred (*errno* is set).

### **Errors:**

EINVAL    The signal *signo* isn't valid.

### **Examples:**

See *sigemptyset()*.



**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*kill()*, *raise()*, *sigaction()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*,  
*sigismember()*, *signal()*, *sigpending()*, *sigprocmask()*

## ***sigblock()***

© 2004, QNX Software Systems Ltd.

*Add to the mask of signals to block*

---

### **Synopsis:**

```
#include <unix.h>

int sigblock(int mask);
```

### **Arguments:**

*mask*     A bitmask of the signals that you want to block.

### **Library:**

`libc`

### **Description:**

The *sigblock()* function adds the signals specified in *mask* to the set of signals currently being blocked from delivery. Signals are blocked if the appropriate bit in *mask* is a 1; the macro *sigmask()* is provided to construct the mask for a given signum. The *sigblock()* returns the previous mask. You can restore the previous mask by calling *sigsetmask()*.

In normal usage, a signal is blocked using *sigblock()*. To begin a critical section, variables modified on the occurrence of the signal are examined to determine that there's no work to be done, and the process pauses awaiting work by using *sigpause()* with the mask returned by *sigblock()*.

It isn't possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

### **Returns:**

The previous set of masked signals.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | No  |

## Caveats:

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multithreaded applications is unsupported.

## See also:

*kill()*, *sigaction()*, *sigmask()*, *signal()*, *sigpause()*, *sigprocmask()*, *sigsetmask()*, *sigunblock()*

## ***sigdelset()***

© 2004, QNX Software Systems Ltd.

*Delete a signal from a set*

### **Synopsis:**

```
#include <signal.h>

int sigdelset(sigset_t *set,
 int signo);
```

### **Arguments:**

*set*        A pointer to the **sigset\_t** object that you want to remove the signal from.

*signo*     The signal that you want to remove. For more information, see “POSIX signals” in the documentation for *SignalAction()*.

### **Library:**

**libc**

### **Description:**

The *sigdelset()* function deletes *signo* from the set pointed to by *set*.

### **Returns:**

0        Success.

-1      An error occurred (*errno* is set).

### **Errors:**

EINVAL    The signal *signo* isn't valid.

### **Examples:**

See *sigemptyset()*.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*kill(), raise(), sigaction(), sigaddset(), sigemptyset(), sigfillset(), sigismember(), signal(), sigpending(), sigprocmask()*

## ***sigemptyset()***

© 2004, QNX Software Systems Ltd.

*Initialize a set to contain no signals*

### **Synopsis:**

```
#include <signal.h>

int sigemptyset(sigset_t *set);
```

### **Arguments:**

*set*     A pointer to the `sigset_t` object that you want to initialize.

### **Library:**

`libc`

### **Description:**

The *sigemptyset()* function initializes *set* to contain no signals.

### **Returns:**

0        Success.

-1       An error occurred (*errno* is set).

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>

void print(sigset_t set, int signo)
{
 printf("Set %8.8lx. Signal %d is ", set, signo);
 if(sigismember(&set, signo))
 printf("a member.\n");
 else
 printf("not a member.\n");
}

int main(void)
{
 sigset_t set;
```

```
sigemptyset(&set);
print(set, SIGINT);

sigfillset(&set);
print(set, SIGINT);

sigdelset(&set, SIGINT);
print(set, SIGINT);

sigaddset(&set, SIGINT);
print(set, SIGINT);
return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*kill()*, *raise()*, *sigaction()*, *sigaddset()*, *sigdelset()*, *sigfillset()*,  
*sigismember()*, *signal()*, *sigpending()*, *sigprocmask()*

# sigevent

© 2004, QNX Software Systems Ltd.

Structure that describes an event

## Synopsis:

```
#include <sys/siginfo.h>

union sigval {
 int sival_int;
 void * sival_ptr;
};
```

The `sigevent` structure is complicated; see below.

## Description:

This structure describes an event. The `int sigev_notify` member indicates how the notification is to occur, as well as which of the other members are used:

| <i>sigev_notify</i>      | <i>sigev_signo</i> | <i>sigev_coid</i> | <i>sigev_priority</i> | <i>sigev_code</i> | <i>sigev_value</i> |
|--------------------------|--------------------|-------------------|-----------------------|-------------------|--------------------|
| SIGEV_INTR               |                    |                   |                       |                   |                    |
| SIGEV_NONE               |                    |                   |                       |                   |                    |
| SIGEV_PULSE              |                    | Connection        | Priority              | Code              | Value              |
| SIGEV_SIGNAL             | Signal             |                   |                       |                   |                    |
| SIGEV_SIGNAL_SIGH        | Signal             |                   |                       | Code              | Value              |
| SIGEV_SIGNAL_SIGR        | Signal             |                   |                       | Code              | Value              |
| SIGEV_THREAD             |                    |                   |                       |                   | Value              |
| (special —<br>see below) |                    |                   |                       |                   |                    |
| SIGEV_UNBLOCK            |                    |                   |                       |                   |                    |

The `<sys/siginfo.h>` file also defines some macros to make initializing the `sigevent` structure easier. All the macros take a pointer to a `sigevent` structure as their first argument, *event*, and set the *sigev\_notify* member to the appropriate value.





---

These macros are QNX Neutrino extensions.

---

### **SIGEV\_INTR**

Raise an interrupt. No other fields in the structure are used.

The initialization macro is:

*SIGEV\_INTR\_INIT( &event )*

### **SIGEV\_NONE**

Don't send any notification. No other fields in the structure are used.

The initialization macro is:

*SIGEV\_NONE\_INIT( &event )*

### **SIGEV\_PULSE**

Send a pulse. The following fields are used:

**int** *sigev\_coid*

The connection ID. This should be attached to the channel with which the pulse will be received.

**short** *sigev\_priority*

The priority of the pulse.

**short** *sigev\_code*

A code to be interpreted by the pulse handler. Although *sigev\_code* can be any 8-bit signed value, you should avoid *sigev\_code* values less than zero in order to avoid conflict with kernel or pulse codes generated by a QNX manager. These codes all start with `_PULSE_CODE_` and are defined in `<sys/neutrino.h>`; for more information, see the documentation for the `_pulse` structure. A safe range of pulse values is `_PULSE_CODE_MINAVAIL` to `_PULSE_CODE_MAXAVAIL`.

**void** \**sigev\_value.sival\_ptr*

A 32-bit value to be interpreted by the pulse handler.

The initialization macro is:

*SIGEV\_PULSE\_INIT( &event, coid, priority, code, value )*

### **SIGEV\_SIGNAL**

Send a signal to a process. The following fields are used:

**int** *sigev\_signo*

The signal to raise. This must be in the range from 1 through NSIG – 1.

The initialization macro is:

*SIGEV\_SIGNAL\_INIT( &event, signal )*

If you need to set the *sigev\_value* for a SIGEV\_SIGNAL event (for example if SA\_SIGINFO is set), you can use this macro:

*SIGEV\_SIGNAL\_VALUE\_INIT( &event, signal, value )*

### **SIGEV\_SIGNAL\_CODE**

This is similar to SIGEV\_SIGNAL, except that SIGEV\_SIGNAL\_CODE also includes a value and a code. The following fields are used:

**int** *sigev\_signo*

The signal to raise. This must be in the range from 1 through NSIG – 1.

**short** *sigev\_code*

A code to be interpreted by the signal handler. This must be in the range from SI\_MINAVAIL through SI\_MAXAVAIL.

**void** \**sigev\_value.sival\_ptr*

A 32-bit value to be interpreted by the signal handler.

The initialization macro is:

*SIGEV\_SIGNAL\_CODE\_INIT( &event, signal, value, code )*

### **SIGEV\_SIGNAL\_THREAD**

Send a signal to a specific thread. The following fields are used:

**int** *sigev\_signo*

The signal to raise. This must be in the range from 1 through NSIG – 1.

**short** *sigev\_code*

A code to be interpreted by the signal handler. This must be in the range from SI\_MINAVAIL through SI\_MAXAVAIL.

**void \****sigev\_value.sival\_ptr*

A 32-bit value to be interpreted by the signal handler.

The initialization macro is:

*SIGEV\_SIGNAL\_THREAD\_INIT( &event, signal, value, code )*

### **SIGEV\_THREAD**

Create a new thread.



---

We don't recommend using this type of event. Pulses are more efficient.

---

The following fields are used:

**void** (*\*sigev\_notify\_function*) (**union sigval**)

A pointer to the function to be notified.

**pthread\_attr \****sigev\_notify\_attributes*

A pointer to thread attributes. This must be NULL, or point to a structure initialized by *pthread\_attr\_init()* at the time of delivery.

`void *sigev_value.sival_ptr`

A value that's to be passed to the notification function.

The initialization macro is:

`SIGEV_THREAD_INIT( &event, fn, value, attr )`

### **SIGEV\_UNBLOCK**

Force a thread to become unblocked. No other fields in the structure are used.

The initialization macro is:

`SIGEV_UNBLOCK_INIT( &event )`

### **Classification:**

QNX Neutrino

### **See also:**

*ds\_create(), InterruptAttach(), InterruptAttachEvent(), iofunc\_notify(), iofunc\_notify\_trigger(), ionotify(), lio\_listio(), mq\_notify(), MsgDeliverEvent(), procmgr\_event\_notify(), \_pulse, TimerCreate(), timer\_create(), TimerInfo(), TimerTimeout(), timer\_timeout()*

“Neutrino IPC” in the Neutrino microkernel chapter of the *System Architecture* guide

**Synopsis:**

```
#include <signal.h>

int sigfillset(sigset_t *set);
```

**Arguments:**

*set*     A pointer to the `sigset_t` object that you want to initialize.

**Library:**

`libc`

**Description:**

The `sigfillset()` function initializes *set* to contain all signals.

**Returns:**

0        Success.  
-1       An error occurred (*errno* is set).

**Examples:**

See `sigemptyset()`.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*kill(), raise(), sigaction(), sigaddset(), sigdelset(), sigemptyset(), sigismember(), signal(), sigpending(), sigprocmask()*

**Synopsis:**

```
#include <signal.h>

int sigismember(const sigset_t *set,
 int signo);
```

**Arguments:**

*set*        A pointer to the **sigset\_t** object that you want to check.

*signo*     The signal that you want to check for membership in the set. For more information, see “POSIX signals” in the documentation for *SignalAction()*.

**Library:**

libc

**Description:**

The *sigismember()* function tests if *signo* is in the set pointed to by *set*.

**Returns:**

1        The *signo* is in the set.

0        The *signo* isn't in the set.

-1      An error occurred (*errno* is set).

**Errors:**

EINVAL    The signal *signo* isn't valid.

## Examples:

See *sigemptyset()*.

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*kill()*, *raise()*, *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*,  
*sigfillset()*, *signal()*, *sigpending()*, *sigprocmask()*



*Restore the environment saved by sigsetjmp(), including the signal mask*

### **Synopsis:**

```
#include <setjmp.h>

void siglongjmp(sigjmp_buf env,
 int val);
```

### **Arguments:**

*env*     The environment saved by the most recent call to *sigsetjmp()*.

*val*     The value that you want *sigsetjmp()* to return.

### **Library:**

`libc`

### **Description:**

The *siglongjmp()* function is a superset of the *longjmp()* function, but also restores the thread's saved signal mask if (and only if) one was saved in the *env* argument by a previous call to *sigsetjmp()*.



---

**WARNING: Don't use *longjmp()* or *siglongjmp()* to restore an environment saved by a call to *setjmp()* or *sigsetjmp()* in another thread. If you're lucky, your application will crash; if not, it'll look as if it works for a while, until random scribbling on the stack causes it to crash.**

---

### **Returns:**

The same values that *longjmp()* returns.

### **Examples:**

See *longjmp()*.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*longjmp()*, *sigaction()*, *sigprocmask()*, *sigsuspend()*

**Synopsis:**

```
#include <unix.h>

#define sigmask(s) (1L<<((s)-1))
```

**Arguments:**

*s*     The signal that you want to create a mask for. For more information, see “POSIX signals” in the documentation for *SignalAction()*.

**Library:**

`libc`

**Description:**

This macro constructs the mask for a given signal number. Use *sigmask()* in conjunction with *sigblock()*, *sigsetmask()*, and *sigunblock()*.

**Returns:**

The signal mask.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **Caveats:**

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multithreaded applications is unsupported.

## **See also:**

*kill()*, *sigaction()*, *sigblock()*, *signal()*, *sigpause()*, *sigsetmask()*, *sigunblock()*

**Synopsis:**

```
#include <signal.h>

void (* signal(int sig,
 void (* func)(int))) (int);
```

**Arguments:**

*sig*      The signal number (defined in `<signal.h>`). For more information, see “POSIX signals” in the documentation for *SignalAction()*.

*func*     The function that you want to call when the signal is raised.

**Library:**

libc

**Description:**

The *signal()* function is used to specify an action to take place when certain conditions are detected while a program executes. See the `<signal.h>` header file for definitions of these conditions, and also refer to the *System Architecture* manual.

There are three types of actions that can be associated with a signal: SIG\_DFL, SIG\_IGN or a pointer to a function. Initially, all signals are set to SIG\_DFL or SIG\_IGN prior to entry of the *main()* routine. An action can be specified for each of the conditions, depending upon the value of the *func* argument, as discussed below.

***func* is a function**

When *func* is a function name, that function is called in a manner equivalent to the following code sequence:

```
/* "sig_no" is condition being signalled */
signal(sig_no, SIG_DFL);
(*func)(sig_no);
```

The *func* function may do the following:

- Return.
- Terminate the program by calling *exit()* or *abort()*.
- Call *longjmp()* or *siglongjmp()*. If you use *longjmp()* to return from a signal handler, the signal remains masked. You can use *siglongjmp()* to restore the mask to the state saved in a previous call to *sigsetjmp()*.

After returning from the signal-catching function, the receiving process resumes execution at the point at which it was interrupted.

The signal catching function is described as follows:

```
void func(int sig_no)
{
 ...
}
```

It isn't possible to catch the SIGSTOP or SIGKILL signals.

Since signal-catching functions are invoked asynchronously with process execution, use the *atomic\_\**, *InterruptLock()*, and *InterruptUnlock()* functions for atomic operations.

### ***func* is SIG\_DFL**

If *func* is SIG\_DFL, the default action for the condition is taken.

If the default action is to stop the process, the execution of that process is temporarily suspended. When a process stops, a SIGCHLD signal is generated for its parent process, unless the parent process has set the SA\_NOCLDSTOP flag (see *sigaction()*). While a process is stopped, any additional signals that are sent to the process aren't delivered until the process is continued, except SIGKILL, which always terminates the receiving process.

Setting a signal action to SIG\_DFL for a signal that is pending, and whose default action is to ignore the signal (for example, SIGCHLD), causes the pending signal to be discarded, whether or not it's blocked.

***func* is SIG\_IGN**

If *func* is SIG\_IGN, the indicated condition is ignored.

You can't set the action for the SIGSTOP and SIGKILL signals to SIG\_IGN.

Setting a signal action to SIG\_IGN for a signal that's pending causes the pending signal to be discarded, whether or not it is blocked.

If a process sets the action for the SIGCHLD signal to SIG\_IGN, the behavior is unspecified.

**Handling a condition**

When a condition is detected, it may be handled by a program, it may be ignored, or it may be handled by the usual default action (often causing an error message to be printed on the *stderr* stream followed by program termination).

A condition can be generated by a program using the *raise()* or *kill()* function

**Returns:**

The previous value of *func* for the indicated condition, or SIG\_ERR if the request couldn't be handled (*errno* is set to EINVAL).

**Examples:**

```
#include <stdlib.h>
#include <signal.h>

sig_atomic_t signal_count;

void MyHandler(int sig_number)
{
 ++signal_count;
}

int main(void)
{
 signal(SIGFPE, MyHandler); /* set own handler */
 signal(SIGABRT, SIG_DFL); /* Default action */
 signal(SIGFPE, SIG_IGN); /* Ignore condition */
 return (EXIT_SUCCESS);
}
```

}

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atomic\_add(), atomic\_add\_value(), atomic\_clr(), atomic\_clr\_value(), atomic\_set(), atomic\_set\_value(), atomic\_sub(), atomic\_sub\_value(), atomic\_toggle(), atomic\_toggle\_value(), InterruptLock(), InterruptUnlock(), kill(), longjmp(), raise(), siglongjmp(), sigprocmask()*



## **SignalAction(), SignalAction\_r()**

*Examine and/or specify actions for signals*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SignalAction(pid_t pid,
 void (* sigstub) (),
 int signo,
 const struct sigaction * act,
 struct sigaction * oact);

int SignalAction_r(pid_t pid,
 void * (sigstub) (),
 int signo,
 const struct sigaction * act,
 struct sigaction * oact);
```

### **Arguments:**

|                |                                                                                                                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>pid</i>     | A process ID, or 0 for the current process.                                                                                                                                                                   |
| <i>sigstub</i> | The address of a signal stub handler. This is a small piece of code in the user's space that interfaces the user's signal handler to the kernel. The library provides a standard one, <i>__signalstub()</i> . |
| <i>signo</i>   | The signal whose action you want to set or get; see "POSIX signals," below.                                                                                                                                   |
| <i>act</i>     | NULL, or a pointer to a <b>sigaction</b> structure that specifies the new action for the signal. For more information, see "Signal actions," below.                                                           |
| <i>oact</i>    | NULL, or a pointer to a <b>sigaction</b> structure where the function can store the old action.                                                                                                               |

### **Library:**

**libc**

## Description:

The *SignalAction()* and *SignalAction\_r()* kernel calls let the calling process examine or specify (or both) the action to be associated with a specific signal in the process *pid*. If *pid* is zero, the calling process is used. The argument *signo* specifies the signal.



---

You should call *sigaction()* instead of using these kernel calls directly.

---

These functions are identical except in the way they indicate errors. See the Returns section for details.

## POSIX signals

The signals are defined in `<signal.h>`, and so are these global variables:

`char * const sys_siglist[]`

An array of signal names.

`int sys_nsig`     The number of entries in the *sys\_siglist* array.

There are 32 POSIX 1003.1a signals, including:

|         |                                              |
|---------|----------------------------------------------|
| SIGHUP  | Hangup.                                      |
| SIGINT  | Interrupt.                                   |
| SIGQUIT | Quit.                                        |
| SIGILL  | Illegal instruction (not reset when caught). |
| SIGTRAP | Trace trap (not reset when caught).          |
| SIGIOT  | IOT instruction.                             |
| SIGABRT | Used by <i>abort()</i> .                     |
| SIGEMT  | EMT instruction.                             |

|          |                                        |
|----------|----------------------------------------|
| SIGFPE   | Floating point exception.              |
| SIGKILL  | Kill (can't be caught or ignored)      |
| SIGBUS   | Bus error.                             |
| SIGSEGV  | Segmentation violation.                |
| SIGSYS   | Bad argument to system call.           |
| SIGPIPE  | Write on pipe with no reader.          |
| SIGALRM  | Realtime alarm clock.                  |
| SIGTERM  | Software termination signal from kill. |
| SIGUSR1  | User-defined signal 1.                 |
| SIGUSR2  | User-defined signal 2.                 |
| SIGCHLD  | Death of child.                        |
| SIGPWR   | Power-fail restart.                    |
| SIGWINCH | Window change.                         |
| SIGURG   | Urgent condition on I/O channel.       |
| SIGPOLL  | System V name for SIGIO.               |
| SIGIO    | Asynchronous I/O.                      |
| SIGSTOP  | Sendable stop signal not from tty.     |
| SIGTSTP  | Stop signal from tty.                  |
| SIGCONT  | Continue a stopped process.            |
| SIGTTIN  | Attempted background tty read.         |
| SIGTTOU  | Attempted background tty write.        |



---

You can't ignore or catch SIGKILL or SIGSTOP.

---

There are 16 POSIX 1003.1b realtime signals, including:

SIGRTMIN      First realtime signal.

SIGRTMAX      Last realtime signal.

The entire range of signals goes from \_SIGMIN (1) to \_SIGMAX (64).

## Signal actions

If *act* isn't NULL, then the specified signal is modified. If *oact* isn't NULL, the previous action is stored in the structure it points to. You can use various combinations of *act* and *oact* to query or set (or both) the action for a signal.

The structure **sigaction** contains the following members:

**void (\*sa\_handler)();**

The address of a signal handler or action for nonqueued signals.

**void (\*sa\_sigaction) (int signo, siginfo\_t \*info, void \*other);**

The address of a signal handler or action for queued signals.

**sigset\_t sa\_mask**

An additional set of signals to be masked (blocked) during execution of the signal-catching function.

**int sa\_flags**

Special flags that affect the behavior of the signal:

- SA\_NOCLDSTOP — don't generate a SIGCHLD on the parent for children who stop via SIGSTOP. This flag is used only when the signal is SIGCHLD.

- SA\_SIGINFO — queue this signal. The default is not to queue a signal delivered to a process. If a signal isn't queued, and the same signal is set multiple times on a process or thread before it runs, only the last signal is delivered. If you set the SA\_SIGINFO flag, the signals are queued, and they're all delivered.

The *sa\_handler* and *sa\_sigaction* members of *act* are implemented as a **union**, and share common storage. They differ only in their prototype, with *sa\_handler* being used for POSIX 1003.1a signals, and *sa\_sigaction* being used for POSIX 1003.1b queued realtime signals. The values stored using either name can be one of:

*function*      The address of a signal-catching function. See below for details.

SIG\_DFL      Use the default action for the signal:

- SIGCHLD, SIGIO, SIGURG, and SIGWINCH — ignore the signal (SIG\_IGN).
- SIGSTOP — stop the process.
- SIGCONT — continue the program.
- All other signals — kill the process.

SIG\_IGN      Ignore the signal. Setting SIG\_IGN for a signal that's pending discards all pending signals, whether it's blocked or not. New signals are discarded. If your process ignores SIGCHLD, its children won't enter the zombie state and the process can't use *wait()* or *waitpid()* to wait on their deaths.

The *function* member of *sa\_handler* or *sa\_sigaction* is always invoked with the following arguments:

```
void handler(int signo, siginfo_t* info, void* other)
```

If you have an old-style signal handler of the form:

```
void handler(int signo)
```

the microkernel passes the extra arguments, but the function simply ignores them.

While in the handler, *signo* is masked, preventing nested signals of the same type. In addition, any signals set in the *sa\_mask* member of *act* are also ORed into the mask. When the handler returns through a normal return, the previous mask is restored, and any pending and now unmasked signals are acted on. You return to the point in the program where it was interrupted. If the thread was blocked in the kernel when the interruption occurred, the kernel call returns with an EINTR (see *ChannelCreate()* and *SyncMutexLock()* for exceptions to this).

When you specify a handler, you must provide the address of a signal stub handler for *sigstub*. This is a small piece of code in the user's space that interfaces the user's signal handler to the kernel. The library provides a standard one, *\_\_signalstub()*.

The *siginfo\_t* structure of the *function* in *sa\_handler* or *sa\_sigaction* contains at least the following members:

- |                            |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>int</b> <i>si_signo</i> | The signal number, which should match the <i>signo</i> argument to the handler.                                                                                                                                                                                                                                                                                                                                                            |
| <b>int</b> <i>si_code</i>  | A signal code, provided by the generator of the signal: <ul style="list-style-type: none"><li>● SI_USER — the <i>kill()</i> function generated the signal.</li><li>● SI_QUEUE — the <i>sigqueue()</i> function generated the signal.</li><li>● SI_TIMER — a timer generated the signal.</li><li>● SI_ASYNCIO — asynchronous I/O generated the signal.</li><li>● SI_MESGQ — POSIX (not QNX) messages queues generated the signal.</li></ul> |

**union sigval** *si\_value*

A value associated with the signal, provided by the generator of the signal.

Signal handlers and actions are defined for the process and affect all threads in the process. For example, if one thread ignores a signal, then all threads ignore the signal.

You can target a signal at a thread, process or process group (see *SignalKill()*). When targeted at a process, at most one thread receives the signal. This thread must have the signal unblocked (see *SignalProcmask()*) to be a candidate for receiving it. All synchronously generated signals (e.g. SIGSEGV) are always delivered to the thread that caused them.

In a multithreaded process, if a signal terminates a thread, by default all threads and thus the process are terminated. You can override this standard POSIX behavior when you create the thread; see *ThreadCreate()*.



---

**CAUTION:** If you use *longjmp()* to return from a signal handler, the signal remains masked. You can use *siglongjmp()* to restore the mask to the state saved by a previous call to *sigsetjmp()*.

---

### Blocking states

These calls don't block.

### Returns:

The only difference between these functions is the way they indicate errors:

*SignalAction()* If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*SignalAction\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

|        |                                                                                                                                                         |
|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The system was unable to allocate a signal handler. This indicated critically low memory.                                                               |
| EFAULT | A fault occurred when the kernel tried to access the buffers provided.                                                                                  |
| EINVAL | The value of <i>signo</i> is less than 1 or greater than <code>_SIGMAX</code> , or you tried to set SIGKILL or SIGSTOP to something other than SIG_DFL. |
| EPERM  | The process doesn't have permission to change the signal actions of the specified process.                                                              |
| ESRCH  | The process indicated by <i>pid</i> doesn't exist.                                                                                                      |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*abort()*, *ChannelCreate()*, *kill()*, *longjmp()*, *siglongjmp()*, *SignalKill()*, *SignalProcmask()*, *sigqueue()*, *sigsetjmp()*, *SyncMutexLock()*, *ThreadCreate()*, *wait()*, *waitpid()*



## ***SignalKill()*, *SignalKill\_r()***

*Send a signal to a process group, process, or thread*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SignalKill(uint32_t nd,
 pid_t pid,
 int tid,
 int signo,
 int code,
 int value);

int SignalKill_r(uint32_t nd,
 pid_t pid,
 int tid,
 int signo,
 int code,
 int value);
```

### **Arguments:**

|                    |                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>nd</i>          | The node descriptor of the node on which to look for <i>pid</i> and <i>tid</i> . To search the local node, set <i>nd</i> to ND_LOCAL_NODE or 0.                                                                                                                                                                        |
| <i>pid</i>         | 0, or the ID of the process to send the signal to; see below.                                                                                                                                                                                                                                                          |
| <i>tid</i>         | 0, or the ID of the thread to send the signal to; see below.                                                                                                                                                                                                                                                           |
| <i>signo</i>       | The signal that you want to send. There are a total of 64 signals available. Of these, at least 8 are POSIX realtime signals that range from SIGRTMIN to SIGRTMAX. For a complete list of signals, see “POSIX signals” in the documentation for <i>SignalAction()</i> . Valid user signals range from 1 to (NSIG - 1). |
| <i>code, value</i> | The code and value associated with the signal; see <i>SignalAction()</i> .                                                                                                                                                                                                                                             |

## Library:

libc

## Description:

The *SignalKill()* and *SignalKill\_r()* kernel calls send the signal *signo* with a code specified by *code* and a value specified by *value* to a process group, process, or thread.

These functions are identical except in the way they indicate errors. See the Returns section for details.

If *signo* is zero, no signal is sent, but the validity of *pid* and *tid* are checked. You can use this as a test for existence.

*SignalKill()* implements the capabilities of the POSIX functions *kill()*, *sigqueue()*, and *pthread\_kill()* in one call. The *pid* and *tid* determine the target of the signal, as follows:

| <i>pid</i> | <i>tid</i> | target                                                             |
|------------|------------|--------------------------------------------------------------------|
| = 0        | —          | Hit the process group of the caller                                |
| < 0        | —          | Hit a process group identified by <i>-pid</i>                      |
| > 0        | = 0        | Hit a single process identified by <i>pid</i>                      |
| > 0        | > 0        | Hit a single thread in process <i>pid</i> identified by <i>tid</i> |

If the target is a thread, the signal is always delivered to exactly that thread. If the thread has the signal blocked — see *SignalProcmask()* — the signal remains pending on the thread.

If the target is a process, the signal is delivered to a thread that has the signal unblocked; see *SignalProcmask()*, *SignalSuspend()*, and *SignalWaitinfo()*. If multiple threads have the signal unblocked, only one thread is given the signal. Which thread receives the signal isn't deterministic. To make it deterministic, you can:

- Have all threads except one block all signals; that thread handles all signals.

Or:

- Target the signals to specific threads.

If all threads have the signal blocked, it's made pending on the process. The first thread to unblock the signal receives the pending signal. If a signal is pending on a thread, it's never retargetted to the process or another thread, regardless of changes to the signal-blocked mask.

If the target is a process group, the signal is delivered as above to each process in the group.

A multithreaded application typically has one thread responsible for catching most or all signals. Threads that don't wish to be directly involved with signals block all signals in their mask.

The signal-blocked mask is maintained on a per-thread basis. The signal-ignore mask and signal handlers are maintained at the process level and are shared by all threads.

If multiple signals are delivered before the target can run and process the signals, the system queues them in priority order if the SA\_SIGINFO bit was set for *signo*. Lower numbered signals have greater priority. If the SA\_SIGINFO bit isn't set for *signo*, then at most one signal is queued at any time. Additional signals with the same *signo* replace existing ones. This is the default behavior for POSIX signal handlers installed using the old *signal()* function. The newer *sigaction()* function lets you control queuing or not on a per-signal basis. Signals with a *code* of SI\_TIMER are never queued.

The *code* and *value* are always saved with the signal. This allows you to deliver data with the signal whether or not SA\_SIGINFO has been set on the *signo*. If SA\_SIGINFO is set, you can use signals to deliver small amounts of data without loss. If you wish to pass significant data, you may wish to consider using *MsgSendPulse()* and *MsgSendv()*, which deliver data with much greater efficiency.

When a thread receives a signal by a signal handler or *SignalWaitinfo()* call, it can retrieve the *signo*, *code* and *value* from a **siginfo\_t** structure, which contains at least the following members:

**int** *si\_signo*      The signal number.  
**int** *si\_code*      The signal code.  
**union sigval** *si\_value*  
                    The signal value.

The value of *si\_code* is limited to an 8-bit signed value as follows:

| Value                       | Description                           |
|-----------------------------|---------------------------------------|
| $-128 \leq si\_code \leq 0$ | User values                           |
| $0 < signo \leq 127$        | System values generated by the kernel |

Some of the common user values defined by POSIX are:

- SLUSER — the *kill()* function generated the signal.
- SLQUEUE — the *sigqueue()* function generated the signal.
- SLTIMER — a timer generated the signal.
- SLASYNCIO — asynchronous I/O generated the signal.
- SLMESGQ — POSIX (not QNX) messages queues generated the signal.

A successful return from this function means the signal has been delivered. What the process(es) or thread does with the signal isn't considered.

If a thread delivers signals that the receiving process has marked as queued faster than the receiver can consume them, the kernel may fail the call if it runs out of signal queue entries. If the *signo*, *code*, and *value* don't change, the kernel performs signal compression by saving an 8-bit count with each queued signal.

**Blocking states**

None. In the network case, lower priority threads may run.

**Returns:**

The only difference between these functions is the way they indicate errors:

|                       |                                                                                                                                                   |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>SignalKill()</i>   | If an error occurs, -1 is returned and sets <i>errno</i> . Any other value returned indicates success.                                            |
| <i>SignalKill_r()</i> | EOK is returned on success. This function does <b>NOT</b> set <i>errno</i> . If an error occurs, any value in the Errors section may be returned. |

**Errors:**

|        |                                                                                                       |
|--------|-------------------------------------------------------------------------------------------------------|
| EINVAL | The value of <i>signo</i> is less than 0 or greater than ( <code>_NSIG - 1</code> ).                  |
| ESRCH  | The process or process group indicated by <i>pid</i> or thread indicated by <i>tid</i> doesn't exist. |
| EPERM  | The process doesn't have permission to send the signal to any receiving process.                      |
| EAGAIN | The kernel had insufficient resources to enqueue the signal.                                          |

**Classification:**

QNX Neutrino

**Safety**

Cancellation point No

Interrupt handler No

*continued...*

## **Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

## **See also:**

*SignalAction(), SignalProcmask(), SignalSuspend(), SignalWaitinfo()*

## **SignalProcmask(), SignalProcmask\_r()**

*Modify or examine the signal-blocked mask of a thread*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SignalProcmask(pid_t pid,
 int tid,
 int how,
 const sigset_t* set,
 sigset_t* oldset);

int SignalProcmask_r(pid_t pid,
 int tid,
 int how,
 const sigset_t* set,
 sigset_t* oldset);
```

### **Arguments:**

- pid*      0, or a process ID; see below.
- tid*      0, or a thread ID; see below.
- how*      The manner in which you want to change the set:
- SIG\_BLOCK — the resulting set is the union of the current set and the signal set pointed to by *set*.
  - SIG\_UNBLOCK — the resulting set is the intersection of the current set and the signal set pointed to by *set*.
  - SIG\_SETMASK — the resulting set is the signal set pointed to by *set*.
- As a special case, you can use the *how* argument to query the current set of pending signals:
- SIG\_PENDING — the combined set of pending signals on the thread and process are saved in the signal set pointed to by *oldset*. The *set* argument is ignored.
- set*      NULL, or a pointer to a **sigset\_t** object that specified the set of signals to be used to change the currently blocked set.

*oldset* NULL, or a pointer to a `sigset_t` object where the function can store the previous blocked mask.

You can use various combinations of *set* and *oldset* to query or change (or both) the signal-blocked mask for a signal.

## Library:

`libc`

## Description:

These kernel calls modify or examine the signal-blocked mask of the thread *tid* in process *pid*. If *pid* is zero, the current process is assumed. If *tid* is zero, *pid* is ignored and the calling thread is used.

The `SignalProcmask()` and `SignalProcmask_r()` functions are identical, except in the way they indicate errors. See the Returns section for details.

When a signal is unmasked, the kernel checks for pending signals on the thread and, if there aren't any pending, checks for pending signals on the process:

| Check                     | Action                                                            |
|---------------------------|-------------------------------------------------------------------|
| Signal pending on thread  | The signal is immediately acted upon.                             |
| Signal pending on process | The signal is moved to the thread and is immediately acted upon.  |
| No signal pending         | No signal action performed until delivery of an unblocked signal. |

It isn't possible to block the SIGKILL or SIGSTOP signals.

When a signal handler is invoked, the signal responsible is automatically masked before its handler is called; see `SignalAction()`. If the handler returns normally, the operating system restores the



signal mask present just before the handler was called as an atomic operation. Changes made using *SignalProcmask()* in the handler are undone.

When a signal is targeted at a process, the kernel delivers it to at most one thread (see *SignalKill()*) that has the signal unblocked. If multiple threads have the signal unblocked, only one thread is given the signal. Which thread receives the signal isn't deterministic. To make it deterministic, you can:

- Have all threads except one block all signals; that thread handles all signals.

Or:

- Target signals to specific threads.

If all threads have the signal blocked, it's made pending on the process. The first thread to unblock the signal receives the pending signal. If a signal is pending on a thread, it's never retargetted to the process or another thread, regardless of changes to the signal-blocked mask.

Signals targeted at a thread always affect that thread alone.

### **Blocking states**

These calls don't block.

### **Returns:**

The only difference between these functions is the way they indicate errors:

#### *SignalProcmask()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

#### *SignalProcmask\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

|        |                                                                                                             |
|--------|-------------------------------------------------------------------------------------------------------------|
| EAGAIN | The system was unable to allocate a signal handler. This indicates critically low memory.                   |
| EFAULT | A fault occurred when the kernel tried to access the buffers provided.                                      |
| EINVAL | The value of <i>how</i> is invalid, or you tried to set SIGKILL or SIGSTOP to something other than SIG_DFL. |
| EPERM  | The process doesn't have permission to change the signal mask of the specified process.                     |
| ESRCH  | The process indicated by <i>pid</i> or thread indicated by <i>tid</i> doesn't exist.                        |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*SignalAction()*, *SignalKill()*

## ***SignalSuspend(), SignalSuspend\_r()***

*Suspend a thread until a signal is received*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SignalSuspend(const sigset_t* set);

int SignalSuspend_r(const sigset_t* set);
```

### **Arguments:**

*set*     A pointer to a **sigset\_t** object that specifies the signals you want to wait for.

### **Library:**

**libc**

### **Description:**

These kernel calls replace the thread's signal mask with the set of signals pointed to by *set* and then suspends the thread until delivery of a signal whose action is either to execute a signal-catching function (then return), or to terminate the thread. On return, the previous signal mask is restored.

The *SignalSuspend()* and *SignalSuspend\_r()* functions are identical, except in the way they indicate errors. See the Returns section for details.

Attempts to block SIGKILL or SIGSTOP are ignored. This is done without causing an error.

If you're using *SignalSuspend()* to synchronously wait for a signal, consider using the more efficient POSIX 1003.1b realtime *sigwaitinfo()* call.

## Blocking states

STATE\_SIGSUSPEND

The calling thread blocks waiting for a signal.

## Returns:

The only difference between these functions is the way they indicate errors.

Since *SignalSuspend()* and *SignalSuspend\_r()* block until interrupted, there's no successful return value.

*SignalSuspend()*

-1 is always returned and *errno* is set.

*SignalSuspend\_r()*

*errno* is **NOT** set, a value in the Errors section is returned.

If the signal handler calls *longjmp()* or *siglongjmp()*, *SignalSuspend()* and *SignalSuspend\_r()* don't return.

## Errors:

|           |                                                                        |
|-----------|------------------------------------------------------------------------|
| EINTR     | The call was interrupted by a signal (this is the normal error).       |
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .       |

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*SignalKill()*

## **SignalWaitinfo(), SignalWaitinfo\_r()** © 2004, QNX Software Systems

Ltd.

*Select a pending signal*

---

### **Synopsis:**

```
#include <sys/neutrino.h>

int SignalWaitinfo(const sigset_t* set,
 siginfo_t* info);

int SignalWaitinfo_r(const sigset_t* set,
 siginfo_t* info);
```

### **Arguments:**

*set*      A pointer to a **sigset\_t** object that specifies the signals you want to wait for.

*info*     NULL, or a pointer to a **siginfo\_t** structure where the function can store information about the signal.

### **Library:**

**libc**

### **Description:**

The *SignalWaitinfo()* and *SignalWaitinfo\_r()* kernel calls select the pending signal from the set specified by *set*. If no signal in *set* is pending at the time of the call, the thread blocks until one or more signals in *set* become pending or until interrupted by an unblocked, caught signal.

These functions are identical except in the way they indicate errors. See the Returns section for details.

If the *info* argument isn't NULL, information on the selected signal is stored there as follows:

| <i>siginfo_t</i> member | Description            |
|-------------------------|------------------------|
| <i>si_signo</i>         | Selected signal number |
| <i>si_code</i>          | Signal code            |
| <i>si_value</i>         | Signal value           |

If, while *SignalWaitinfo()* is waiting, a caught signal occurs that isn't blocked, the signal handler is invoked and *SignalWaitinfo()* is interrupted with an error of EINTR.

### Blocking states

STATE\_SIGWAITINFO

The calling thread blocks waiting for a signal.

### Returns:

The only difference between these functions is the way they indicate errors:

*SignalWaitinfo()*

A signal number. If an error occurs, -1 is returned and *errno* is set.

*SignalWaitinfo\_r()*

A signal number. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

### Errors:

|           |                                                                        |
|-----------|------------------------------------------------------------------------|
| EINTR     | The call was interrupted by a signal.                                  |
| EFAULT    | A fault occurred when the kernel tried to access the buffers provided. |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .       |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*SignalKill()*, *SignalKill\_r()*



## ***significand()*, *significandf()***

*Compute the “significant bits” of a floating-point number*

### **Synopsis:**

```
#include <math.h>

double significand (double x);

float significandf (float x);
```

### **Arguments:**

*x*     A floating-point number.

### **Library:**

`libm`

### **Description:**

The *significand()* and *significandf()* functions are math functions that compute the “significant bits” of a floating-point number.

When encoding a floating-point number into binary notation, you remove the sign, and then shift the bits to the right or left until the shifted result is in the range [0.5, 1). The negative of the number of positions shifted is the *exponent* of the number, and the shifted result is the *significand*.

If  $x$  equals  $\text{sig} \times 2^n$  with  $1 < \text{sig} < 2$ , then `significand(x)` returns `sig` for exercising the fraction-part(F) test vector. The function `significand(x)` isn't defined when  $x$  is one of:

- 0
- positive or negative infinity
- NAN.

## Returns:

*scalb* ( *x*, (double) -ilogb ( *x* ) )

Since *significand*(*x*) = *scalb*(*x*, -ilogb(*x*)) where *ilogb*() returns the exponent part of *x* and *scalb*(*x*, *n*) returns *a*, such that  $x = a \times 2^n$ , then:

**When *x* is: *scalbn*(*x*, *n*) returns:**

|           |          |
|-----------|----------|
| ±infinity | <i>x</i> |
| NAN       | NAN      |

## Examples:

```
#include <stdio.h>
#include <errno.h>
#include <inttypes.h>
#include <math.h>
#include <fpstatus.h>

int main(int argc, char** argv)
{
 double a, b, c, d;

 a = 5;
 b = ilogb(d);
 printf("The exponent part of %f is %f \n", a, b);
 c = significand(a);
 printf("%f = %f * (2 ^ %f) \n", a, c, b);

 return(0);
}
```

produces the output:

```
The exponent part of 5.000000 is -895.000000
5.000000 = 1.250000 * (2 ^ -895.000000)
```

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ilogb()*, *scalb()*, *scalbn()*

# ***sigpause()***

© 2004, QNX Software Systems Ltd.

*Wait for a signal*

---

## **Synopsis:**

```
#include <signal.h>

int sigpause(int sig);
```

## **Arguments:**

*sig*     A mask containing the signal number that you want to wait for.

## **Library:**

`libc`

## **Description:**

The *sigpause()* function assigns *sig* to the set of masked signals and then waits for a signal to arrive; on return, the set of masked signals is restored. The *mask* argument is usually 0 to indicate that no signals are now to be blocked. This function always terminates by being interrupted, returning -1, and setting *errno* to EINTR.

In normal usage, a signal is blocked using *sigblock()*. To begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause()* with the mask returned by *sigblock()*.

It isn't possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

## **Returns:**

-1; *errno* is set to EINTR.

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multithreaded applications is unsupported.

## See also:

*kill()*, *sigaction()*, *sigblock()*, *sigmask()*, *signal()*, *sigsetmask()*, *sigsuspend()*, *sigunblock()*

## ***sigpending()***

© 2004, QNX Software Systems Ltd.

*Examine the set of pending, masked signals for a process*

### **Synopsis:**

```
#include <signal.h>

int sigpending(sigset_t *set);
```

### **Arguments:**

*set* A pointer to a **sigset\_t** object that the function sets to indicate the pending, masked signals.

### **Library:**

libc

### **Description:**

The *sigpending()* function is used to examine the set of pending signals that are masked (blocked) from delivery to the calling thread and that are pending on the calling process or thread. They're saved in the signal set pointed to by *set*.

### **Returns:**

0 Success.  
-1 An error occurred (*errno* is set).

### **Errors:**

EFAULT A fault occurred while accessing the buffer pointed to by *set*.

### **Examples:**

See *sigprocmask()*.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*kill(), raise(), sigaction(), sigaddset(), sigdelset(), sigemptyset(), sigfillset(), sigismember(), signal(), sigprocmask()*

## ***sigprocmask()***

© 2004, QNX Software Systems Ltd.

*Examine or change the signal mask for a thread*

### **Synopsis:**

```
#include <signal.h>

int sigprocmask(int how,
 const sigset_t *set,
 sigset_t *oset);
```

### **Arguments:**

- how*     The manner in which you want to change the set:
- SIG\_BLOCK — add the signals pointed to by *set* to the thread mask.
  - SIG\_UNBLOCK — remove the signals pointed to by *set* from the thread mask.
  - SIG\_SETMASK — set the thread mask to be the signals pointed to by *set*.
- set*     NULL, or a pointer to a **sigset\_t** object that defines the signals that you want to change in the thread's signal mask. If this argument is NULL, the *how* argument is ignored.
- oset*    NULL, or a pointer to a **sigset\_t** object that the function sets to indicate the thread's current signal mask.

### **Library:**

**libc**

### **Description:**

The *sigprocmask()* function is used to examine or change (or both) the signal mask for the calling thread. If the value of *set* isn't NULL, it points to a set of signals to be used to change the currently blocked set.

The *set* argument isn't changed. The resulting set is maintained in the process table of the calling thread. If a signal occurs on a signal that's



masked, it becomes pending, but doesn't affect the execution of the process. You can examine pending signals by calling *sigpending()*. When a pending signal is unmasked, it's acted upon immediately, before this function returns.

When a signal handler is invoked, the signal responsible is automatically masked before its handler is called. If the handler returns normally, the operating system restores the signal mask present just before the handler was called as an atomic operation. Changes made using *sigprocmask()* in the handler are undone.

The *sigaction()* function lets you specify any mask that's applied before a handler is invoked. This can simplify multiple signal handler design.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EAGAIN Insufficient system resources are available to mask the signals.
- EFAULT A fault occurred trying to access the buffers provided.
- EINVAL The value of *how* is invalid, or you tried to set SIGKILL or SIGSTOP to something other than SIG\_DFL.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <signal.h>
#include <unistd.h>

int main(void)
{
 sigset_t set, oset, pset;

 sigemptyset(&set);
```

```
sigaddset(&set, SIGINT);
sigprocmask(SIG_BLOCK, &set, &oset);
printf("Old set was %8.8ld.\n", oset);

sigpending(&pset);
printf("Pending set is %8.8ld.\n", pset);

kill(getpid(), SIGINT);

sigpending(&pset);
printf("Pending set is %8.8ld.\n", pset);

sigprocmask(SIG_UNBLOCK, &set, &oset);

/* The program terminates with a SIGINT */
return(EXIT_SUCCESS);
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*kill()*, *raise()*, *sigaction()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *sigismember()*, *signal()*, *SignalProcmask()*, *sigpending()*

## Synopsis:

```
#include <signal.h>

int sigqueue (pid_t pid,
 int signo,
 const union sigval value);
```

## Arguments:

- pid*        The ID of the process that you want to signal.
- signo*      Zero, or the number of the signal that you want to queue for the process. For more information, see “POSIX signals” in the documentation for *SignalAction()*.
- value*      The value to queue with the signal.

## Library:

`libc`

## Description:

The *sigqueue()* function causes the signal, *signo* to be sent with the specified value to the process, *pid*. If *signo* is zero, error checking is performed, but no signal is actually sent. This is one way of checking to see if *pid* is valid.

The condition required for a process to have permission to queue a signal to another process is the same as for the *kill()* function — the real or effective user ID of the sending process must match the real or effective user ID of the receiving process.

The *sigqueue()* function returns immediately. If SA\_SIGINFO is set for *signo* and if the resources are available to queue the signal, the signal is queued and sent to the receiving process. If SA\_SIGINFO isn't set for the *signo*, then *signo* is sent to the receiving process if the signal isn't already pending.

If *pid* causes *signo* to be generated for the sending process, and if *signo* isn't blocked for the calling thread and if no other thread has *signo* unblocked or is waiting in a *sigwait()* function for *signo*, then either *signo* or at least one pending unblocked signal is delivered to the calling thread before *sigqueue()* returns.

Should any of multiple pending signals in the range SIGRTMIN to SIGRTMAX be selected for delivery, the lowest numbered one is delivered. The selection order between realtime and nonrealtime signals, or between multiple pending nonrealtime signals, is unspecified.

## Returns:

- 0 Success.
- 1 An error occurred; *errno* is set.

## Errors:

- EAGAIN No resources were available to queue the signal. The process has already queued the maximum number of signals as returned by:  
*sysconf( \_SC\_SIGQUEUE\_MAX )*  
that are still pending at the receiver(s), or a system-wide resource limit has been exceeded.
- EINVAL The value of the *signo* argument is an invalid or unsupported signal number.
- ENOSYS The function *sigqueue()* isn't supported by this implementation.
- EPERM The process doesn't have the appropriate privilege to send the signal to the receiving process.
- ESRCH The process *pid* doesn't exist.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*kill()*, *signal()*

## ***sigsetjmp()***

© 2004, QNX Software Systems Ltd.

*Save the environment, including the signal mask*

### **Synopsis:**

```
#include <setjmp.h>

int sigsetjmp(sigjmp_buf env,
 int savemask);
```

### **Arguments:**

|                 |                                                                             |
|-----------------|-----------------------------------------------------------------------------|
| <i>env</i>      | A buffer where the function can save the calling environment.               |
| <i>savemask</i> | Nonzero if you want to save the process's current signal mask, otherwise 0. |

### **Library:**

libc

### **Description:**

The *sigsetjmp()* function behaves in the same way as the *setjmp()* function when *savemask* is zero. If *savemask* is nonzero, then *sigsetjmp()* also saves the thread's current signal mask as part of the calling environment.



---

**WARNING:** *Don't use `longjmp()` or `siglongjmp()` to restore an environment saved by a call to `setjmp()` or `sigsetjmp()` in another thread. If you're lucky, your application will crash; if not, it'll look as if it works for a while, until random scribbling on the stack causes it to crash.*

---

### **Returns:**

Zero on the first call, or nonzero if the return is the result of a call to *siglongjmp()*.

**Examples:**

See *setjmp()*.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*setjmp()*, *sigaction()*, *sigprocmask()*, *sigsuspend()*

## ***sigsetmask()***

© 2004, QNX Software Systems Ltd.

*Set the mask of signals to block*

---

### **Synopsis:**

```
#include <unix.h>

int sigsetmask(int mask);
```

### **Arguments:**

*mask*     A bitmask of the signals that you want to block.

### **Library:**

`libc`

### **Description:**

The *sigsetmask()* function sets the current signal mask (those signals that are blocked from delivery). Signals are blocked if the corresponding bit in *mask* is a 1; the macro *sigmask()* is provided to construct the mask for a given signal.

In normal usage, a signal is blocked using *sigblock()*. To begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause()* with the mask returned by *sigblock()*.

It isn't possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

### **Returns:**

The previous set of masked signals.

### **Classification:**

Unix



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multithreaded applications is unsupported.

**See also:**

*kill()*, *sigaction()*, *sigblock()*, *sigmask()*, *signal()*, *sigpause()*, *sigprocmask()*, *sigunblock()*

## ***sigsuspend()***

© 2004, QNX Software Systems Ltd.

*Replace the signal mask, and then suspend the thread*

### **Synopsis:**

```
#include <signal.h>

int sigsuspend(const sigset_t *sigmask);
```

### **Arguments:**

*sigmask*     A pointer to a **sigset\_t** object that specifies the signals that you want in the thread's signal mask.

### **Library:**

libc

### **Description:**

The *sigsuspend()* function replaces the thread's signal mask with the set of signals pointed to by *sigmask* and then suspends the thread until delivery of a signal whose action is either to execute a signal-catching function (then return), or to terminate the thread.

### **Returns:**

-1 (if the function returns); *errno* is set.

### **Errors:**

EFAULT     A fault occurred trying to access the buffers provided.

EINTR     A signal was caught by the calling thread, and control is returned from the signal-catching function.

### **Examples:**

```
/*
 * This program pauses until a signal other than
 * a SIGINT occurs. In this case a SIGALRM.
 */
#include <stdio.h>
#include <signal.h>
```

```
#include <stdlib.h>
#include <unistd.h>

sigset_t set;

int main(void)
{
 sigemptyset(&set);
 sigaddset(&set, SIGINT);

 printf("Program suspended and immune to breaks.\n");
 printf("A SIGALRM will terminate the program"
 " in 10 seconds.\n");
 alarm(10);
 sigsuspend(&set);
 return(EXIT_SUCCESS);
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pause()*, *sigaction()*, *sigpending()*, *sigprocmask()*

## ***sigtimedwait()***

© 2004, QNX Software Systems Ltd.

*Wait for a signal or a timeout*

### **Synopsis:**

```
#include <signal.h>

int sigtimedwait(const sigset_t *set,
 siginfo_t *info,
 const struct timespec *timeout);
```

### **Arguments:**

|                |                                                                                                                                                                                                                                                 |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>set</i>     | A set of signals from which the function selects a pending signal.                                                                                                                                                                              |
| <i>info</i>    | If this argument is NULL, the selected signal is returned by <i>sigwaitinfo()</i> ; otherwise, the selected signal is stored in the <i>si_signo</i> member of <i>info</i> , and the cause of the signal is stored in the <i>si_code</i> member. |
| <i>timeout</i> | NULL, or a pointer to a <b>timespec</b> structure that specifies the maximum time to wait for a pending signal.                                                                                                                                 |

### **Library:**

**libc**

### **Description:**

The *sigtimedwait()* function selects a pending signal from *set*, atomically clears it from the set of pending signals in the process, and returns that signal number.

If any value is queued to the selected signal, the first queued value is dequeued and, if the *info* argument is non-NULL, the value is stored in the *si\_value* member of *info*. The system resources used to queue the signal are released and made available to queue other signals. If no value is queued, the content of the *si\_value* member is undefined.

If no further signals are queued for the selected signal, the pending indication for that signal is reset.

If none of the signals specified by *set* are pending, *sigtimedwait()* waits for the time interval specified by the **timespec** structure *timeout*. If *timeout* is zero and if none of the signals specified by *set* are pending, then *sigtimedwait()* returns immediately with an error. If *timeout* is NULL, *sigtimedwait()* behaves the same as *sigwaitinfo()*.

**Returns:**

The selected signal number, or -1 if an error occurred (*errno* is set).

**Errors:**

|        |                                                                                                                                                                                      |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The <i>timeout</i> expired before a signal specified in <i>set</i> was generated, or all kernel timers are in use.                                                                   |
| EFAULT | A fault occurred while accessing the provided buffers.                                                                                                                               |
| EINTR  | The wait was interrupted by an unblocked, caught signal.                                                                                                                             |
| EINVAL | The <i>timeout</i> argument specified a <i>tv_nsec</i> value less than zero or greater than or equal to 1000 million or <i>set</i> contains an invalid or unsupported signal number. |

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pause()*, *pthread\_sigmask()*, *sigaction()*, *SignalWaitinfo()*,  
*sigpending()*, *sigsuspend()*, *sigwaitinfo()*, **timespec**

## Synopsis:

```
#include <unix.h>

int sigunblock(int mask);
```

## Arguments:

*mask*     A bitmask of the signals that you want to unblock.

## Library:

`libc`

## Description:

The *sigunblock()* function removes the signals specified in *mask* from the set of signals currently being blocked from delivery. Signals are unblocked if the appropriate bit in *mask* is a 1; the macro *sigmask()* is provided to construct the mask for a given signal. The *sigunblock()* returns the previous mask. You can restore the previous mask by calling *sigsetmask()*.

In normal usage, a signal is blocked using *sigblock()*. To begin a critical section, variables modified on the occurrence of the signal are examined to determine that there is no work to be done, and the process pauses awaiting work by using *sigpause()* with the mask returned by *sigblock()*.

It isn't possible to block SIGKILL, SIGSTOP, or SIGCONT; this restriction is silently imposed by the system.

## Returns:

The previous set of masked signals.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

Use of these interfaces should be restricted to only applications written on BSD platforms. Use of these interfaces with any of the system libraries or in multithreaded applications is unsupported.

**See also:**

*kill(), sigaction(), sigblock(), sigmask(), signal(), sigpause(), sigprocmask(), sigsetmask()*



## Synopsis:

```
#include <signal.h>

int sigwait(const sigset_t *set,
 int *sig);
```

## Arguments:

- set*     A pointer to a **sigset\_t** object that specifies the signals you want to wait for.
- sig*     A pointer to a location where the function can store the signal that it cleared.

## Library:

**libc**

## Description:

The *sigwait()* function selects a pending signal from *set*, atomically clears it from the set of pending signals in the system, and returns that signal number in *sig*. If there are multiple signals queued for the signal number selected, the first signal causes a return from *sigwait()* and the rest remain queued. If no signal in *set* is pending at the time of the call, the thread is suspended until one or more becomes pending.

The signals defined by *set* must be blocked before you call *sigwait()*; otherwise, the behavior is undefined. The effect of *sigwait()* on the signal actions for the signals in *set* is unspecified.

If more than one thread is using *sigwait()* to wait for the same signal, only one of the threads returns from *sigwait()* with the signal number — which one is unspecified.

**Returns:**

- 0 Success (that is, one of the signals specified by set is pending or has been generated).
- EINTR The *sigwait()* function was interrupted by a signal.
- EINVAL The *set* argument contains an invalid or unsupported signal number.
- EFAULT A fault occurred while accessing the provided buffers.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pause()*, *pthread\_sigmask()*, *sigaction()*, *sigpending()*, *sigsuspend()*

### **Synopsis:**

```
#include <signal.h>

int sigwaitinfo (const sigset_t *set,
 siginfo_t *info);
```

### **Arguments:**

*set*      A pointer to a **sigset\_t** object that specifies the signals you want to wait for.

*info*     NULL, or a pointer to a **siginfo\_t** structure where the function can store information about the signal.

### **Library:**

**libc**

### **Description:**

The *sigwaitinfo()* function selects a pending signal from *set*, atomically clears it from the set of pending signals in the system, and returns that signal number.

If *info* is NULL, the selected signal is returned by *sigwaitinfo()*; otherwise, the selected signal is stored in the *si\_signo* member of *info* and the cause of the signal is stored in the *si\_code* member.

If any value is queued to the selected signal, the first queued value is dequeued and, if the *info* argument is non-NULL, the value is stored in the *si\_value* member of *info*. The system resources used to queue the signal are released and made available to queue other signals. If no value is queued, the content of the *si\_value* member is undefined.

If no further signals are queued for the selected signal, the pending indication for that signal is reset.

**Returns:**

A signal number, or -1 if an error occurred (*errno* is set).

**Errors:**

EFAULT     A fault occurred while accessing the buffers.

EINTR     The wait was interrupted by an unblocked, caught signal.

**Classification:**

POSIX 1003.1

**Safety**

---

Cancellation point    Yes

Interrupt handler     No

Signal handler        Yes

Thread                 Yes

**See also:**

*pause()*, *pthread\_sigmask()*, *sigaction()*, *SignalWaitinfo()*,  
*sigpending()*, *sigsuspend()*, *sigtimedwait()*, *sigwait()*

## Synopsis:

```
#include <math.h>

double sin(double x);

float sinf(float x);
```

## Arguments:

*x* The angle, in radians, for which you want to compute the sine.

## Library:

`libm`

## Description:

The *sin()* and *sinf()* functions compute the sine (specified in radians) of *x*. An argument with a large magnitude may yield a result with little or no significance.

## Returns:

The sine value.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main(void)
{
 printf("%f\n", sin(.5));
}
```

```
 return(EXIT_SUCCESS);
}
```

produces the output:

0.479426

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*acos()*, *asin()*, *atan()*, *atan2()*, *cos()*, *tan()*

## Synopsis:

```
#include <math.h>

double sinh(double x);

float sinhf(float x);
```

## Arguments:

*x* The angle, in radians, for which you want to compute the hyperbolic sine.

## Library:

libm

## Description:

The *sinh()* and *sinhf()* functions compute the hyperbolic sine (specified in radians) of *x*. A range error occurs if the magnitude of *x* is too large.

## Returns:

The hyperbolic sine value.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main(void)
{
 printf("%f\n", sinh(.5));
}
```

```
 return(EXIT_SUCCESS);
}
```

produces the output:

0.521095

### Classification:

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*cosh()*, *errno*, *tanh()*



## Synopsis:

```
#include <unistd.h>

unsigned int sleep(unsigned int seconds);
```

## Arguments:

*seconds*     The number of realtime seconds that you want to suspend the thread for.

## Library:

libc

## Description:

The *sleep()* function suspends the calling thread until the number of realtime seconds specified by the *seconds* argument have elapsed, or the thread receives a signal whose action is either to terminate the process or to call a signal handler. The suspension time may be greater than the requested amount, due to the scheduling of other, higher priority threads by the system.

## Returns:

0 if the full time specified was completed; otherwise, the number of seconds unslept if interrupted by a signal.

## Errors:

EAGAIN     No timer resources were available to satisfy the request.

## Examples:

```
/*
 * The following program sleeps for the
 * number of seconds specified in argv[1].
 */
#include <stdlib.h>
#include <unistd.h>
```

```
int main(int argc, char **argv)
{
 unsigned seconds;

 seconds = (unsigned) strtol(argv[1], NULL, 0);
 sleep(seconds);

 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*alarm()*, *delay()*, *errno*, *nanosleep()*, *timer\_create()*, *timer\_gettime()*,  
*timer\_settime()*, *usleep()*

### **Synopsis:**

```
#include <pthread.h>

int _sleepon_broadcast(sleepon_t * l,
 const volatile void * addr);
```

### **Arguments:**

- l*            A pointer to a **sleepon\_t** that you created by calling *\_sleepon\_init()*.
- addr*        The handle that the threads are waiting on. The value of *addr* is typically a data structure that controls a resource.

### **Library:**

**libc**

### **Description:**

The *\_sleepon\_signal()* and *\_sleepon\_broadcast()* functions are very similar:

- *\_sleepon\_signal()* wakes up a single thread that's waiting on the key, *addr*.
- *\_sleepon\_broadcast()* wakes up all threads that are waiting on the key, *addr*.

The waiting threads must have used the same *sleepon*, *l* and key, *addr*, in order to be woken up.

To be woken up, the calling threads must have been locked by *\_sleepon\_lock()*.

### **Returns:**

- 0            Success.
- ≠0         Failure.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*pthread\_sleepon\_broadcast(), pthread\_sleepon\_signal(),  
\_sleepon\_destroy(), \_sleepon\_init(), \_sleepon\_lock(), \_sleepon\_signal(),  
\_sleepon\_unlock()*

## Synopsis:

```
#include <pthread.h>

int _sleepon_destroy(sleepon_t * l);
```

## Arguments:

*l* A pointer to a **sleepon\_t** that you created by calling *\_sleepon\_init()*.

## Library:

libc

## Description:

The *\_sleepon\_destroy()* function destroys a **sleepon\_t** structure, *l*, that has been previously initialized by *\_sleepon\_init()*.

If *l* hasn't been locked by *\_sleepon\_lock()*, *\_sleepon\_destroy()* locks it before destroying it.

The sleepon structure is reference-counted such that, if other threads are blocked waiting for a condition, they're be signaled to wake up, and the last one to wake up frees the memory allocated to the sleepon.

## Returns:

0 Success.  
≠0 Failure

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*\_sleepon\_broadcast()*, *\_sleepon\_init()*, *\_sleepon\_lock()*,  
*\_sleepon\_signal()*, *\_sleepon\_unlock()*

## Synopsis:

```
#include <pthread.h>

int _sleepon_init(sleepon_t ** pl,
 unsigned flags);
```

## Arguments:

*pl*            The address of a location where the function can store a pointer to the **sleepon\_t** object that it creates.

*flags*        There are currently no flags defined; pass zero for this argument.

## Library:

libc

## Description:

The *\_sleepon\_init()* function allocates a **sleepon\_t** object (which is an opaque data structure) and stores a pointer to it in the location that *pl* points to.

## Returns:

0            Success.

≠0         Failure.

## Classification:

QNX Neutrino

### **Safety**

---

Cancellation point   No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*\_sleepon\_broadcast()*, *\_sleepon\_destroy()*, *\_sleepon\_lock()*,  
*\_sleepon\_signal()*, *\_sleepon\_unlock()*



**Synopsis:**

```
#include <pthread.h>

int _sleepon_lock(sleepon_t * l);
```

**Arguments:**

*l* A pointer to a **sleepon\_t** that you created by calling *\_sleepon\_init()*.

**Library:**

libc

**Description:**

The *\_sleepon\_lock()* function locks the mutex associated with the sleepon structure, *l*.

You must call this function before calling *\_sleepon\_wait()*, *\_sleepon\_signal()*, or *\_sleepon\_broadcast()*.

**Returns:**

|         |                                                                                                |
|---------|------------------------------------------------------------------------------------------------|
| EOK     | Success.                                                                                       |
| EAGAIN  | Insufficient system resources were available to lock the mutex.                                |
| EDEADLK | The calling thread already owns <i>mutex</i> , and the mutex doesn't allow recursive behavior. |
| EINVAL  | Invalid mutex.                                                                                 |

The *\_sleepon\_lock()* function returns the same values as *pthread\_mutex\_lock()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_lock(), \_sleepon\_broadcast(), \_sleepon\_destroy(),  
\_sleepon\_init(), \_sleepon\_signal(), \_sleepon\_unlock(), \_sleepon\_wait()*

## Synopsis:

```
#include <pthread.h>

int _sleepon_signal(sleepon_t * l,
 const volatile void * addr);
```

## Arguments:

- l*            A pointer to a **sleepon\_t** that you created by calling *\_sleepon\_init()*.
- addr*        The handle that the thread is waiting on. The value of *addr* is typically a data structure that controls a resource.

## Library:

**libc**

## Description:

The *\_sleepon\_signal()* and *\_sleepon\_broadcast()* functions are very similar:

- *\_sleepon\_signal()* wakes up a single thread that's waiting on the key, *addr*.
- *\_sleepon\_broadcast()* wakes up all threads that are waiting on the key, *addr*.

The waiting threads must have used the same *sleepon*, *l* and key, *addr*, in order to be woken up.

To be woken up, the calling threads must have been locked by *\_sleepon\_lock()*.

## Returns:

- 0            Success.
- ≠0         Failure.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*pthread\_sleepon\_broadcast()*, *pthread\_sleepon\_signal()*,  
*\_sleepon\_broadcast()*, *\_sleepon\_destroy()*, *\_sleepon\_init()*,  
*\_sleepon\_lock()*, *\_sleepon\_unlock()*

## Synopsis:

```
#include <pthread.h>

int __sleepon_unlock(sleepon_t * l);
```

## Arguments:

*l* A pointer to a `sleepon_t` that you created by calling `__sleepon_init()`.

## Library:

`libc`

## Description:

The `__sleepon_unlock()` function unlocks the mutex associated with the sleepon structure, *l*. You must have previously locked the mutex by calling `__sleepon_lock()`.

## Returns:

EOK Success.

EINVAL Invalid mutex *mutex*.

EPERM The current thread doesn't own *mutex*.

The `__sleepon_unlock()` function returns the same values as `pthread_mutex_unlock()`.

## Classification:

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*pthread\_mutex\_unlock()*, *\_sleepon\_broadcast()*, *\_sleepon\_destroy()*,  
*\_sleepon\_init()*, *\_sleepon\_lock()*

## Synopsis:

```
#include <pthread.h>

int _sleepon_wait(sleepon_t * l,
 const volatile void * addr,
 _uint64 nsec);
```

## Arguments:

- l* A pointer to a **sleepon\_t** that you created by calling *\_sleepon\_init()*.
- addr* The handle that you want to wait on. The value of *addr* is typically a data structure that controls a resource.
- nsec* Zero, or the amount of time, in nanoseconds, to wait before the thread wakes up. If this timeout occurs, ETIMEDOUT is returned.

## Library:

**libc**

## Description:

The *\_sleepon\_wait()* function blocks on the sleepon *l* using the key *addr* until woken up by either a *\_sleepon\_signal()* or a *\_sleepon\_broadcast()* call that uses the same *addr* key.

The calling thread must first have locked the sleepon by calling *\_sleepon\_lock()*.

When the thread returns from this function, it must release the sleepon lock by calling *\_sleepon\_unlock()*.

## Returns:

- 0 Success.
- ≠0 Failure; a nonzero *errno* value.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*\_sleepon\_broadcast()*, *\_sleepon\_lock()*, *\_sleepon\_unlock()*,  
*\_sleepon\_signal()*



## Synopsis:

```
#include <stdio.h>
#include <sys/slog.h>

int slogb(int code,
 int severity,
 void * data,
 int size);
```

## Arguments:

|                 |                                                                                                   |
|-----------------|---------------------------------------------------------------------------------------------------|
| <i>opcode</i>   | A combination of a <i>major</i> and <i>minor</i> code.                                            |
| <i>severity</i> | The severity of the log message; see “Severity levels,” in the documentation for <i>slogf()</i> . |
| <i>data</i>     | A block of raw data.                                                                              |
| <i>size</i>     | The size of the raw data.                                                                         |

## Library:

`libc`

## Description:

The *slog\*()* functions send log messages to the system logger, **slogger**. To send formatted messages, use *slogf()*. If you have programs that scan log files for specified codes, you can use *slogb()* or *slogi()* to send a block of structures or `int`'s, respectively.

## Errors:

Any value from the Errors section in *MsgSend()*, as well as:

|        |                                                                                           |
|--------|-------------------------------------------------------------------------------------------|
| EACCES | Insufficient permission to write to the log file.                                         |
| EINVAL | The size of the data buffer exceeds 255×4 bytes, or an odd number of bytes is being read. |

ENOENT      Invalid log file or directory specified, or **slogger** isn't running.

**Examples:**

See *slogf()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*slogf()*, *slogi()*, *vslogf()*

**slogger**, **sloginfo** in the *Utilities Reference*

**Synopsis:**

```
#include <stdio.h>
#include <sys/slog.h>

int slogf(int opcode,
 int severity,
 const char * fmt,
 ...);
```

**Arguments:**

- opcode*      A combination of a *major* and *minor* code. Create the *opcode* using the `_SLOG_SETCODE(major, minor)` macro that's defined in `<sys/slog.h>`.  
The *major* and *minor* codes are defined in `<sys/slogcodes.h>`.
- severity*     The severity of the log message; see "Severity levels," below.
- fmt*         A standard `printf()` string followed by `printf()` arguments.

The formatting characters that you use in the message determine any additional arguments.

**Library:**

`libc`

**Description:**

The `slog*()` functions send log messages to the system logger, `slogger`. To send formatted messages, use `slogf()`. If you have programs that scan log files for specified codes, you can use `slogb()` or `slogi()` to send a block of structures or `int`'s, respectively.

The `vslogf()` function is an alternate form in which the arguments have already been captured using the variable-length argument facilities of `<stdarg.h>`.

## Severity levels

There are eight levels of severity defined. The lowest severity is 7 and the highest is 0. The default is 7.

| <b>Manifest Name</b> | <b>Severity value</b> | <b>Description</b>                                                        |
|----------------------|-----------------------|---------------------------------------------------------------------------|
| ._SLOG_SHUTDOWN      | 0                     | Shut down the system NOW (e.g. for OEM use)                               |
| ._SLOG_CRITICAL      | 1                     | Unexpected unrecoverable error (e.g. hard disk error)                     |
| ._SLOG_ERROR         | 2                     | Unexpected recoverable error (e.g. needed to reset a hardware controller) |
| ._SLOG_WARNING       | 3                     | Expected error (e.g. parity error on a serial port)                       |
| ._SLOG_NOTICE        | 4                     | Warnings (e.g. out of paper)                                              |
| ._SLOG_INFO          | 5                     | Information (e.g. printing page 3)                                        |
| ._SLOG_DEBUG1        | 6                     | Debug messages (normal detail)                                            |
| ._SLOG_DEBUG2        | 7                     | Debug messages (fine detail)                                              |

## Returns:

The size of the message sent to **slogger**, or -1 if an error occurs.

## Errors:

Any value from the Errors section in *MsgSend()*, as well as:

|        |                                                                                           |
|--------|-------------------------------------------------------------------------------------------|
| EACCES | Insufficient permission to write to the log file.                                         |
| EINVAL | The size of the data buffer exceeds 255×4 bytes, or an odd number of bytes is being read. |
| ENOENT | Invalid log file or directory specified, or <b>slogger</b> isn't running.                 |

## Examples:

```
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>
#include <sys/slog.h>
#include <sys/slogcodes.h>

int main() {
 int i;

 for(i = 0 ; ; i++) {
 switch(rand() % 3) {
 case 0:
 slogb(_SLOG_SETCODE(_SLOGC_TEST, 0),
 _SLOG_DEBUG1, &i, sizeof(i));
 break;

 case 1:
 slogi(_SLOG_SETCODE(_SLOGC_TEST, 1),
 _SLOG_CRITICAL, 1, i);
 break;

 case 2:
 slogf(_SLOG_SETCODE(_SLOGC_TEST, 2),
 _SLOG_ERROR,
 "This is number %d", i);
 break;
 }

 sleep(1);
 }
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*slogb()*, *slogi()*, *vslogf()*

**slogger**, **sloginfo** in the *Utilities Reference*

**Synopsis:**

```
#include <stdio.h>
#include <sys/slog.h>

int slogi(int code,
 int severity,
 int nargs,
 ...);
```

**Arguments:**

*opcode*      A combination of a *major* and *minor* code.

*severity*     The severity of the log message; see “Severity levels,” in the documentation for *slogf()*.

*nargs*        The number of integers to send. A maximum of 32 integers is allowed.

The additional arguments are the integers that you want to write.

**Library:**

`libc`

**Description:**

The *slog\*()* functions send log messages to the system logger, **slogger**. To send formatted messages, use *slogf()*. If you have programs that scan log files for specified codes, you can use *slogb()* or *slogi()* to send a block of structures or **int**'s, respectively.

**Errors:**

Any value from the Errors section in *MsgSend()*, as well as:

EACCES      Insufficient permission to write to the log file.

EINVAL      The size of the data buffer exceeded 32 integers.

ENOENT Invalid log file or directory specified, or **slogger** isn't running.

**Examples:**

See *slogf()*

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*slogb()*, *slogf()*, *slogi()*, *vslogf()*

**slogger**, **sloginfo** in the *Utilities Reference*



## Synopsis:

```
#include <malloc.h>

void* _smalloc(size_t size);
```

## Arguments:

*size*     The size of the block to allocate, in bytes.

## Library:

`libc`

## Description:

The `_smalloc()` function allocates space for an object of *size* bytes. Nothing is allocated when the *size* argument has a value of zero.



---

This function allocates memory in blocks of `_amblksiz` bytes; `_amblksiz` is a global variable defined in `<stdlib.h>`.

---

You must use `_sfree()` to deallocate the memory allocated by `_smalloc()`.

## Returns:

A pointer to the start of the allocated memory, or NULL if there's insufficient memory available, or if the requested *size* is zero.

## Classification:

QNX Neutrino

### Safety

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | No  |
| Thread            | Yes |

**See also:**

*calloc()*, *free()*, *realloc()*, *\_scalloc()*, *\_sfree()*, *\_srealloc()*

## Synopsis:

```
#include <sys/types.h>
#include <snmp/asn1.h>
#include <snmp/snmp_api.h>

extern int snmp_errno

int snmp_close(struct snmp_session * session);
```

## Arguments:

*session*     A pointer to the `snmp_session` structure that identifies the SNMP session that you want to close. This pointer was returned by a call to `snmp_open()`.

## Library:

`libsnmp`

## Description:

The `snmp_close()` function closes the input session, frees any data allocated for it, dequeues any pending requests, and closes any sockets allocated for the session.

## Returns:

- 1     Success.
- 0     An error occurred (*snmp\_errno* is set).

## Errors:

If an error occurs, this function sets *snmp\_errno* to:

SNMPERR\_BAD\_SESSION

The specified session wasn't open.

## Classification:

SNMP

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*select()*, *snmp\_free\_pdu()*, *snmp\_open()*, **snmp\_pdu**,  
*snmp\_pdu\_create()*, *snmp\_read()*, *snmp\_select\_info()*, *snmp\_send()*,  
**snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

**Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <snmp/asn1.h>
#include <snmp/snmp_api.h>

void snmp_free_pdu(struct snmp_pdu * pdu);
```

**Arguments:**

*pdu* A pointer to the `snmp_pdu` structure that you want to free.

**Library:**

`libsnmp`

**Description:**

The `snmp_free_pdu()` function frees the `snmp_pdu` structure pointed to by *pdu*, and any data that it contains that was allocated with `malloc()`.

**Classification:**

SNMP

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

### **See also:**

*select()*, *snmp\_close()*, *snmp\_open()*, **snmp\_pdu**, *snmp\_read()*,  
*snmp\_select\_info()*, *snmp\_send()*, **snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

## Synopsis:

```
#include <sys/types.h>
#include <snmp/asn1.h>
#include <snmp/snmp_api.h>

extern int snmp_errno;

struct snmp_session * snmp_open(
 struct snmp_session * session);
```

## Arguments:

*session*     A pointer to a **snmp\_session** structure that defines the SNMP session that you want to open.

## Library:

libsnmp

## Description:

The *snmp\_open()* function sets up an SNMP session with the information supplied by the application in the **snmp\_session** structure pointed to by *session*. Next, *snmp\_open()* opens and binds the necessary UDP port.

## Returns:

A pointer to a **snmp\_session** structure for the created session (which is different from the pointer passed to the function), or NULL if an error occurs (*snmp\_errno* is set).

## Errors:

If an error occurs, this function sets *snmp\_errno* to one of:

SNMPERR\_BAD\_ADDRESS  
Unknown host.

SNMPERR.BAD.LOCPORT

Couldn't bind to the specified port.

SNMPERR.GENERR

Couldn't open the socket.

## Classification:

SNMP

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*read\_main\_config\_file()*, *select()*, *snmp\_close()*, *snmp\_free\_pdu()*,  
**snmp\_pdu**, *snmp\_pdu\_create()*, *snmp\_read()*, *snmp\_select\_info()*,  
*snmp\_send()*, **snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)



**Synopsis:**

```
#include <snmp/snmp_api.h>

struct snmp_pdu {
 int version;
 ipaddr address;
 oid * srcParty;
 int srcPartyLen;
 oid * dstParty;
 int dstPartyLen;
 oid * context;
 int contextLen;
 u_char * community;
 int community_len;
 int command;
 long reqid;
 long errstat;
 long errindex;

 /* Trap information */
 oid * enterprise;
 int enterprise_length;
 ipaddr agent_addr;
 int trap_type;
 int specific_type;
 u_long time;

 struct variable_list * variables;
};
```

**Description:**

The `snmp_pdu` structure describes a *Protocol Data Unit* (PDU), a transaction that's performed over an open session. It contains the headers and variables of an SNMP packet. The structure includes the following members:

|                |                                                                                          |
|----------------|------------------------------------------------------------------------------------------|
| <i>version</i> | The version of SNMP: either <code>SNMP_VERSION_1</code> or <code>SNMP_VERSION_2</code> . |
|----------------|------------------------------------------------------------------------------------------|

|                         |                                                                                                                                         |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <i>address</i>          | The destination IP address.                                                                                                             |
| <i>srcParty</i>         | The source party being used.                                                                                                            |
| <i>srcPartyLen</i>      | The number of object identifier (OID) elements in <i>srcParty</i> . For example, if <i>srcParty</i> is <b>.1.3.6</b> , the length is 3. |
| <i>dstParty</i>         | The destination party being used.                                                                                                       |
| <i>dstPartyLen</i>      | The number of OID elements in <i>dstParty</i> .                                                                                         |
| <i>context</i>          | The context being used.                                                                                                                 |
| <i>contextLen</i>       | The number of OID elements in <i>context</i> .                                                                                          |
| <i>community</i>        | The community for outgoing requests.                                                                                                    |
| <i>communityLen</i>     | The length of the community name.                                                                                                       |
| <i>command</i>          | The type of this PDU.                                                                                                                   |
| <i>reqid</i>            | The request ID. The default is <code>SNMP_DEFAULT_REQID (0)</code> .                                                                    |
| <i>errstat</i>          | The error status ( <code>non_repeaters</code> in <code>GetBulk</code> ). The default is <code>SNMP_DEFAULT_ERRSTAT (-1)</code> .        |
| <i>errindex</i>         | The error index ( <code>max_repetitions</code> in <code>GetBulk</code> ). The default is <code>SNMP_DEFAULT_ERRINDEX (-1)</code> .      |
| <i>enterprise</i>       | The system OID.                                                                                                                         |
| <i>enterpriseLength</i> | The number of OID elements in <i>enterprise</i> . The default is <code>SNMP_DEFAULT_ENTERPRISE_LENGTH (0)</code> .                      |
| <i>agent_addr</i>       | The address of the object generating the trap.                                                                                          |
| <i>trap_type</i>        | The trap type.                                                                                                                          |
| <i>specific_type</i>    | The specific type.                                                                                                                      |

- time*            The up time. The default is SNMP\_DEFAULT\_TIME (0).
- variables*        A linked list of variables, of type `variable_list`.

The `variable_list` structure is defined as:

```
typedef struct sockaddr_in ipaddr;

struct variable_list {
 struct variable_list* next_variable;
 oid* name;
 int name_length;
 u_char type;
 union {
 long* integer;
 u_char* string;
 oid* objid;
 u_char* bitstring;
 struct counter64* counter64;
 } val;
 int val_len;
};
```

The members are:

- next\_variable*    A pointer to the next variable. This is NULL for the last variable in the list.
- name*            The object identifier of the variable.
- name\_length*     The number of sub IDs in *name*.
- type*            The ASN type of variable.
- val.integer*     The value of the variable if it's an integer.
- val.string*      The value of the variable if it's a string.
- val.objid*       The value of the variable if it's an object ID.
- bitstring*       The value of the variable if it's a bitstring.
- counter64*       The value of the variable if it's a counter64.
- val\_len*         The length of the value.

**Classification:**

SNMP

**See also:**

*snmp\_close()*, *snmp\_free\_pdu()*, *snmp\_open()*, *snmp\_pdu\_create()*,  
*snmp\_read()*, *snmp\_send()*, **snmp\_session**

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

### **Synopsis:**

```
#include <sys/types.h>
#include <netinet/in.h>
#include <snmp/snmp.h>
#include <snmp/asn1.h>
#include <snmp/snmp_api.h>
#include <snmp/snmp_client.h>

extern int snmp_errno;

struct snmp_pdu * snmp_pdu_create (int command);
```

### **Arguments:**

*command* The type of message that the PDU represents:

- BULK\_REQ\_MSG
- GET\_REQ\_MSG
- GET\_RSP\_MSG
- GETNEXT\_REQ\_MSG
- INFORM\_REQ\_MSG
- SET\_REQ\_MSG
- TRP\_REQ\_MSG
- TRP2\_REQ\_MSG

as defined in `<snmp/snmp.h>`.

### **Library:**

libsnmp

### **Description:**

The *snmp\_pdu\_create()* function allocates memory for a *Protocol Data Unit* (PDU) structure for SNMP message passing. The PDU structure is initialized with default values; see the documentation for `snmp_pdu`.

## Returns:

A pointer to the PDU structure created, or NULL if an error occurs (*snmp\_errno* is set).

## Errors:

If an error occurs, this function sets *snmp\_errno* to:

SNMPERR\_GENERR

Not enough memory to create the *snmp\_pdu* structure.

## Classification:

SNMP

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*snmp\_free\_pdu()*, *snmp\_pdu*, *snmp\_read()*, *snmp\_send()*

*RFC 1157*, FAQ in Internet newsgroup `comp.protocols.snmp`

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

## Synopsis:

```
#include <sys/select.h>
#include <snmp/snmp_impl.h>

void snmp_read(struct fd_set * fdset);
```

## Arguments:

*fdset*     A pointer to a **fd\_set** structure that contains all the file descriptors that you want to read from.

## Library:

**libsnmp**

## Description:

The *snmp\_read()* function reads a packet from each socket and its set of file descriptors and parses the packet. The resulting Protocol Data Unit (PDU) is passed to the callback routine for the session (see **snmp\_session**); if the callback returns successfully, the PDU and its request are deleted.

For information on asynchronous SNMP transactions, see *snmp\_select\_info()*.

## Classification:

SNMP

### Safety

Cancellation point    Yes

Interrupt handler      No

Signal handler        No

Thread                 No

## **See also:**

*select()*, *snmp\_close()*, *snmp\_open()*, **snmp\_pdu**, *snmp\_read()*,  
*snmp\_select\_info()*, *snmp\_send()*, **snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)



**Synopsis:**

```
#include <sys/types.h>
#include <sys/select.h>
#include <sys/time.h>
#include <snmp/snmp_api.h>

int snmp_select_info(int * numfds,
 struct fd_set * fdset,
 struct timeval * timeout,
 int * block);
```

**Arguments:**

- numfds*      The number of significant file descriptors in *fdset*.
- fdset*        A pointer to a set of file descriptors that contains all of the file descriptors that you've opened for SNMP. If activity occurs on any of these file descriptors, you should call *snmp\_read()* with that file-descriptor set.
- timeout*     A pointer to a **timeval** structure that defines the longest time that SNMP can wait for a timeout. You should call *select()* with the minimum time between *timeout* and any other timeouts necessary. You should check this on each invocation of *select()*. If a timeout is received, you should call *snmp\_timeout()* to see if the timeout was for SNMP. (The *snmp\_timeout()* function is idempotent.)
- You must provide the *timeout*, even if *block* is 1 (see below).
- block*        Governs the behavior of *select()*:
- If *block* is 0, *select()* is requested to time out.
  - If *block* is 1, *select()* is requested to block indefinitely. The timeout value is treated as undefined, although you must provide it. On return, if *block* is nonzero, the value of *timeout* is undefined.

## Library:

libsnmp

## Description:

The *snmp\_select\_info()* function is used to return information about what SNMP requires from a *select()* call.

## Asynchronous SNMP transactions:

To have SNMP transactions occur asynchronously, you can invoke the functions *snmp\_select\_info()*, *snmp\_timeout()*, and *snmp\_read()* in conjunction with the system call *select()*. For more information, see *select()*.

For asynchronous transactions, invoke *snmp\_select\_info()* with the information you would have passed to *select()* in the absence of SNMP. The *snmp\_select\_info()* function modifies the information, which is subsequently passed to *select()*.

| Parameters to <i>select()</i> : | Corresponding parameters to <i>snmp_select_info()</i> : |
|---------------------------------|---------------------------------------------------------|
|---------------------------------|---------------------------------------------------------|

*nfds*

*numfds*

*readfds*

*fdset*

*timeout*

*timeout* — must point to an allocated (but not necessarily initialized) **timeval** structure.

The following code segment shows how to use these SNMP functions in conjunction with *select()*:

```
FD_ZERO (&fdset);
numfds=sd+1;
FD_SET (sd, &fdset);
block=0;
tvp=&timeout;
timerclear (tvp);
tvp->tv_sec = 5;
```

```
snmp_select_info(&numfds, &fdset, tvp, &block);

if(block==1)
{
 tvp = NULL;
}
count = select(numfds, &fdset, 0, 0, tvp);

if(count==0)
 snmp_timeout();
if(count>0)
 snmp_read(&fdset);
```

**Returns:**

The number of open sockets (i.e. the number of open sessions).

**Classification:**

SNMP

**Safety**

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**See also:**

*select()*, *snmp\_close()*, *snmp\_open()*, **snmp\_pdu**, *snmp\_read()*,  
*snmp\_select\_info()*, *snmp\_send()*, **snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

## ***snmp\_send()***

© 2004, QNX Software Systems Ltd.

*Send SNMP messages*

### **Synopsis:**

```
#include <sys/types.h>
#include <snmp/asn1.h>
#include <netinet/in.h>
#include <snmp/snmp_api.h>

extern int snmp_errno

int snmp_send(struct snmp_session * session,
 struct snmp_pdu * pdu);
```

### **Arguments:**

*session*     A pointer to the **snmp\_session** structure that identifies the SNMP session that you want to send the message on. This pointer was returned by a call to *snmp\_open()*.

*pdu*         A pointer to the **snmp\_pdu** structure that defines the Protocol Data Unit that you want to send. Create this structure by calling *snmp\_pdu\_create()*.

### **Library:**

libsnmp

### **Description:**

The *snmp\_send()* function sends the PDU on the session provided. If necessary, some of the **snmp\_pdu** structure data is set from the session defaults. A request corresponding to this PDU is added to the list of outstanding requests on this session and then the packet is sent.

This function frees *pdu* unless an error occurs.

### **Returns:**

The request ID of the generated packet, if applicable, 1 if not applicable, or 0 if an error occurs (*snmp\_errno* is set).

**Errors:**

If an error occurs, this function sets *snmp\_errno* to one of:

**SNMPERR\_BAD\_ADDRESS**

A necessary entity in the *pdu* structure was omitted. These include:

- *version*
- *address* and the `snmp_session.peername` member
- *srcParty* (SNMP version 2 only)
- *dstParty* (SNMP version 2 only)
- *context* (SNMP version 2 only)
- *communityLen* (SNMP version 1 only)

**SNMPERR\_BAD\_SESSION**

The specified session wasn't open.

**SNMPERR\_GENERR**

An error occurred forming the packet.

**Classification:**

SNMP

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

### **See also:**

*select()*, *snmp\_close()*, *snmp\_open()*, **snmp\_pdu**, *snmp\_pdu\_create()*,  
*snmp\_read()*, *snmp\_select\_info()*, *snmp\_send()*, **snmp\_session**,  
*snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

**Synopsis:**

```
#include <snmp/snmp_api.h>

struct snmp_session {
 u_char * community;
 int community_len;
 int retries;
 long timeout;
 char * peername;
 u_short remote_port;
 u_short local_port;
 u_char * (*authenticator) ();
 int (* callback) ();
 void * callback_magic;
 int version;
 oid * srcParty;
 int srcPartyLen;
 oid * dstParty;
 int dstPartyLen;
 oid * context;
 int contextLen;
};
```

**Description:**

The **snmp\_session** structure describes a set of transactions sharing similar transport characteristics. It includes the following members:

|                      |                                                                                                                                           |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <i>community</i>     | The community for outgoing requests. The default is 0.                                                                                    |
| <i>community_len</i> | The length of the community name. The default is SNMP_DEFAULT_COMMUNITY_LEN (0).                                                          |
| <i>retries</i>       | The number of retries before timing out. The default is SNMP_DEFAULT_RETRIES (-1).                                                        |
| <i>timeout</i>       | The number of microseconds until the first timeout. Subsequent timeouts increase exponentially. The default is SNMP_DEFAULT_TIMEOUT (-1). |

|                      |                                                                                                                                                                                                      |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>peername</i>      | The domain name or dotted IP address of the default peer. The default is <code>SNMP_DEFAULT_PEERNAME</code> (NULL).                                                                                  |
| <i>remote_port</i>   | The UDP port number of the peer. The default is <code>SNMP_DEFAULT_REMPORT</code> (0).                                                                                                               |
| <i>local_port</i>    | My UDP port number. The default is <code>SNMP_DEFAULT_ADDRESS</code> (0), for picked randomly.                                                                                                       |
| <i>authenticator</i> | The authentication function, or NULL if null authentication is used. If your application is using version 1 of SNMP, you must supply this member. The <i>authenticator()</i> function is defined as: |

```
u_char* authenticator(u_char* pdu,
 int* length,
 u_char* community,
 int community_len)
```

The arguments are:

- *pdu* — the rest of the PDU to be authenticated.
- *length* — the length of the remaining data in the PDU, updated by *authenticator()*.
- *community* — the community name for authentication.
- *community\_len* — the length of the community name.

To specify null authentication, set the *authenticator* field in `snmp_session` to NULL.

The *authenticator()* function returns an authenticated PDU, or NULL if an error occurs.

|                 |                                                                                                             |
|-----------------|-------------------------------------------------------------------------------------------------------------|
| <i>callback</i> | A function used to extract the data from the received packet (the <code>snmp_pdu</code> structure passed to |
|-----------------|-------------------------------------------------------------------------------------------------------------|



the callback). The application must supply this member.

The *callback()* function is defined as:

```
int callback(int operation,
 struct snmp_session* session,
 int reqid,
 struct snmp_pdu* pdu,
 void* magic);
```

The arguments are:

- *operation* — the possible operations are RECEIVED\_MESSAGE and TIMED\_OUT.
- *session* — the session that was authenticated using *community*.
- *reqid* — the request ID identifying the transaction within this session. Use 0 for traps.
- *pdu* — A pointer to PDU information. You must copy the information because it will be freed elsewhere.
- *magic* — a pointer to the data for *callback()*.

The callback should return 1 on successful completion, or 0 if it should be kept pending.

|                       |                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| <i>callback_magic</i> | A pointer to data that the callback function may consider important.                                                                    |
| <i>version</i>        | The version of SNMP: either SNMP_VERSION_1 or SNMP_VERSION_2.                                                                           |
| <i>srcParty</i>       | The source party being used for this session.                                                                                           |
| <i>srcPartyLen</i>    | The number of object identifier (OID) elements in <i>srcParty</i> . For example, if <i>srcParty</i> is <i>.1.3.6</i> , the length is 3. |
| <i>dstParty</i>       | The destination party being used for this session.                                                                                      |

*dstPartyLen*      The number of OID elements in *dstParty*.  
*context*            The context being used for this session.  
*contextLen*        The number of OID elements in *context*.

### Classification:

SNMP

### See also:

*snmp\_close()*, *snmp\_free\_pdu()*, *snmp\_open()*, **snmp\_pdu**,  
*snmp\_send()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

## Synopsis:

```
#include <snmp/snmp_api.h>

void snmp_timeout(void);
```

## Library:

libsnmp

## Description:

The *snmp\_timeout()* function handles any outstanding SNMP requests. It should be called whenever the timeout from *snmp\_select\_info()* expires. The *snmp\_timeout()* function checks to see if any of the sessions has an outstanding request that has timed out.

If it finds one or more, and that PDU has more retries available, a new packet is formed from the PDU and is resent. If there are no more retries available, the callback for the session is used to alert the user of the timeout by setting the callback's *operation* argument to TIMED\_OUT (2).

For information on asynchronous SNMP transactions, see *snmp\_select\_info()*.

## Classification:

SNMP

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

### **See also:**

*select()*, *snmp\_close()*, *snmp\_open()*, **snmp\_pdu**, *snmp\_read()*,  
*snmp\_select\_info()*, *snmp\_send()*, **snmp\_session**, *snmp\_timeout()*

*RFC 1157*, FAQ in Internet newsgroup **comp.protocols.snmp**

Marshall T. Rose, *The Simple Book: An Introduction to Internet Management*, Revised 2nd ed. (Prentice-Hall, 1996, ISBN 0-13-451659-1)

*Write formatted output to a character array, up to a given maximum number of characters*

### **Synopsis:**

```
#include <stdio.h>

int snprintf(char* buf,
 size_t count,
 const char* format,
 ...);
```

### **Arguments:**

- buf* A pointer to the buffer where you want the function to store the formatted string.
- count* The maximum number of characters to store in the buffer, including a terminating null character.
- format* A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

### **Library:**

`libc`

### **Description:**

The *snprintf()* function is similar to *fprintf()*, except that *snprintf()* places the generated output into the character array pointed to by *buf*, instead of writing it to a file. A null character is placed at the end of the generated character string.

### **Returns:**

The number of characters that would have been written into the array, not counting the terminating null character, had *count* been large enough. It does this even if *count* is zero; in this case *buf* can be NULL.

If an error occurred, *snprintf()* returns a negative value and sets *errno*.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

/* Create temporary file names using a counter */

char namebuf[13];
int TempCount = 0;

char *make_temp_name(void)
{
 snprintf(namebuf, 13, "ZZ%.6o.TMP",
 TempCount++);
 return(namebuf);
}

int main(void)
{
 FILE *tf1, *tf2;

 tf1 = fopen(make_temp_name(), "w");
 tf2 = fopen(make_temp_name(), "w");
 fputs("temp file 1", tf1);
 fputs("temp file 2", tf2);
 fclose(tf1);
 fclose(tf2);

 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

Be careful if you're using *snprintf()* to build a string one piece at a time. For example, this code:

```
len += snprintf(&buf[len], RECSIZE - 1 - len, ...);
```

could have a problem if *snprintf()* truncates the string. Without a separate test to compare *len* with `RECSIZE`, this code doesn't protect against a buffer overflow. After the call that truncates the output, *len* is larger than `RECSIZE`, and `RECSIZE - 1 - len` is a very large (unsigned) number; the next call generates unlimited output somewhere beyond the buffer.

## See also:

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *sprintf()*, *swprintf()*, *vfprintf()*, *vfwprintf()*, *vprintf()*, *vsprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*, *wprintf()*

## ***socketmark()***

© 2004, QNX Software Systems Ltd.

*Determine whether a socket is at the out-of-band mark*

### **Synopsis:**

```
#include <sys/socket.h>

int socketmark(int s);
```

### **Arguments:**

*s*     The file descriptor of the socket that you want to check, as returned by *socket()*.

### **Library:**

`libsocket`

### **Description:**

The *socketmark()* function determines whether the socket specified by *s* is at the out-of-band data mark. If the protocol for the socket supports out-of-band data by marking the stream with an out-of-band data mark, *socketmark()* returns 1 when all data preceding the mark has been read and the out-of-band data mark is the first element in the receive queue.

The *socketmark()* function doesn't remove the out-of-band data mark from the stream.

Using this function between receive operations lets an application determine which data comes before and after out-of-band data.

### **Returns:**

0     The socket isn't at the out-of-band data mark.  
1     The socket is at the out-of-band data mark.  
-1    An error occurred (*errno* is set).



## Errors:

- EBADF      Invalid file descriptor *s*.
- ENOTTY     The *s* argument isn't the file descriptor of a valid socket.

## Classification:

POSIX 1003.1-2001

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*recv()*, *recvmsg()*

# socket()

© 2004, QNX Software Systems Ltd.

Create an endpoint for communication

## Synopsis:

```
#include <sys/types.h>
#include <sys/socket.h>

int socket(int domain,
 int type,
 int protocol);
```

## Arguments:

*domain* The communications domain that you want to use. This selects the protocol family that should be used. These families are defined in `<sys/socket.h>`.

*type* The type of socket you want to create. This determines the semantics of communication. Here are the currently defined types:

- SOCK\_STREAM — provides sequenced, reliable, two-way, connection-based byte streams. An out-of-band data transmission mechanism may be supported.
- SOCK\_DGRAM — supports datagrams, which are connectionless, unreliable messages of a fixed (typically small) maximum length.
- SOCK\_RAW — provides access to internal network protocols and interfaces. Available only to the superuser, this type isn't described here.

For more information, see below.

*protocol* The particular protocol that you want to use with the socket. Normally, only a single protocol exists to support a particular socket type within a given protocol family. But if many protocols exist, you must specify one. The protocol number you give is particular to the communication domain where communication is to take place (see `/etc/protocols` in the *Utilities Reference*).

## Library:

`libsocket`

## Description:

The *socket()* function creates an endpoint for communication and returns a descriptor.

### SOCK\_STREAM sockets

SOCK\_STREAM sockets are full-duplex byte streams, similar to pipes. A stream socket must be in a *connected* state before any data may be sent or received on it. To create a connection to another socket, call *connect()* call.

Once the socket is connected, you can transfer data by using *read()* and *write()* or some variant of *send()* and *recv()*. When a session has been completed, a *close()* may be performed. Out-of-band data may also be transmitted (as described in *send()*) and received (as described in *recv()*).

The communications protocols used to implement a SOCK\_STREAM socket ensure that data isn't lost or duplicated. If a piece of data that the peer protocol has buffer space for can't be successfully transmitted within a reasonable length of time, the connection is considered broken and calls will indicate an error by returning -1 and setting *errno* to ETIMEDOUT.

### SOCK\_DGRAM and SOCK\_RAW sockets

With SOCK\_DGRAM and SOCK\_RAW sockets, datagrams can be sent to correspondents named in *send()* calls. Datagrams are generally received with *recvfrom()*, which returns the next datagram with its return address.

### Using the *ioctl()* call

You can use the *ioctl()* call to specify a process group to receive a SIGURG signal when the out-of-band data arrives. The call may also enable nonblocking I/O and asynchronous notification of I/O events via SIGIO.

**Socket-level options**

The operation of sockets is controlled by socket-level options. These options are defined in the file `<sys/socket.h>`. Use `setsockopt()` and `getsockopt()` to set and get options.

**Returns:**

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

**Errors:**

|                 |                                                                                                        |
|-----------------|--------------------------------------------------------------------------------------------------------|
| EACCES          | Permission to create a socket of the specified type and/or protocol is denied.                         |
| EMFILE          | The per-process descriptor table is full.                                                              |
| ENFILE          | The system file table is full.                                                                         |
| ENOBUFS         | Insufficient buffer space available. The socket can't be created until sufficient resources are freed. |
| ENOMEM          | Not enough memory.                                                                                     |
| EPROTONOSUPPORT | The protocol type or the specified protocol isn't supported within this domain.                        |

**Classification:**

Standard Unix, POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## Caveats:

By default, *socket()* communicates with the TCP/IP stack managing the `/dev/socket` portion of the namespace. This behavior can be controlled via the `SOCK` environmental variable. See the examples in the `npm-tcpip.so` utility.

## See also:

ICMP6, ICMP, INET6, IPv6, IP, IPsec, ROUTE, TCP, UDP, UNIX protocols

*accept()*, *bind()*, *close()*, *connect()*, *getprotobyname()*, *getsockname()*, *getsockopt()*, *ioctl()*, *listen()*, *read()*, *recv()*, *select()*, *send()*, *shutdown()*, *socketpair()*, *write()*

## ***socketpair()***

© 2004, QNX Software Systems Ltd.

*Create a pair of connected sockets*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/socket.h>

int socketpair(int domain,
 int type,
 int protocol,
 int * fd[2]);
```

### **Arguments:**

|                 |                                                                                                                                                                    |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>domain</i>   | The communications domain where the sockets are to be created.                                                                                                     |
| <i>type</i>     | The type of sockets to create.                                                                                                                                     |
| <i>protocol</i> | The protocol to use with the sockets. A protocol of 0 causes <i>socketpair()</i> to use an unspecified default protocol appropriate for the requested socket type. |
| <i>fd[2]</i>    | The 2-digit integer array where the file descriptors of the created socket pair are to be held.                                                                    |

### **Library:**

**libsocket**

### **Description:**

The *socketpair()* call creates an unnamed pair of connected sockets in the specified *domain*, of the specified *type*, using the optionally specified *protocol* argument. The file descriptors are returned in the vector *fd* and are identical.

Valid types are described in *socket()*.

If the *protocol* argument is nonzero, it must be a protocol that's understood by the address family. No such protocols are defined at this time.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EAFNOSUPPORT  
The specified address family isn't supported on this machine.
- EFAULT  
The address *sv* doesn't specify a valid part of the process address space.
- EMFILE  
Too many descriptors are in use by this process.
- EOPNOTSUPP  
The specified protocol doesn't support creation of socket pairs.
- EPROTONOSUPPORT  
The specified protocol isn't supported on this machine.

**Examples:**

```
#include <stdio.h>
#include <sys/socket.h>

#define CHAR_BUFSIZE 20
int main(int argc, char **argv) {
 int fd[2], len;
 char message[CHAR_BUFSIZE];

 if(socketpair(AF_LOCAL, SOCK_STREAM, 0, fd) == -1) {
 return 1;
 }

 /* Print a message into one end of the socket */
 snprintf(message, CHAR_BUFSIZE, "First message");
 write(fd[0], message, strlen(message) + 1);

 /* Print a message into the other end of the socket */
 snprintf(message, CHAR_BUFSIZE, "Second message");
```

```
write(fd[1], message, strlen(message) + 1);

/* Read back the data written to the first socket */
len = read(fd[0], message, CHAR_BUF_SIZE-1);
message[len] = '\0';
printf("Read [%s] from first fd \n", message);

/* Read back the data written to the second socket */
len = read(fd[1], message, CHAR_BUF_SIZE-1);
message[len] = '\0';
printf("Read [%s] from second fd \n", message);

close(fd[0]);
close(fd[1]);

return 0;
}
```

**Classification:**

POSIX 1003.1-2001

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:***read(), socket(), write()*



### **Synopsis:**

```
#include <sys/select.h>

int SOCKSinit(char * progname);
```

### **Arguments:**

*progname*      The name that you want to associate with the your program. The default is **SOCKSclient**.

### **Library:**

**libsocks**

### **Description:**

The *SOCKSinit()* function initializes some defaults for the SOCKS library and also sets the program name that appears in the **syslog** output.

You don't have to call this function before making a SOCKS library call (but if you don't, a generic "SOCKSclient" appears instead of the program name).

For more information about SOCKS and its libraries, see the appendix, SOCKS — A Basic Firewall.

### **Returns:**

- 0      Success.
- 1      An error occurred (*errno* is set).

### **Classification:**

SOCKS

## Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*Raccept()*, *Rbind()*, *Rconnect()*, *Rgetsockname()*, *Rlisten()*, *Rrcmd()*,  
*Rselect()*

SOCKS — A Basic Firewall

**Synopsis:**

```
#include <unistd.h>
#include <fcntl.h>
#include <sys/stat.h>
#include <sys/types.h>
#include <share.h>

int sopen(const char* filename,
 int oflag,
 int share,
 ...);
```

**Arguments:**

- filename*     The path name of the file that you want to open.
- oflag*        Flags that specify the status and access modes of the file. This argument is a combination of the following bits (defined in <fcntl.h>):
- O\_RDONLY — permit the file to be only read.
  - O\_WRONLY — permit the file to be only written.
  - O\_RDWR — permit the file to be both read and written.
  - O\_APPEND — cause each record that's written to be written at the end of the file.
  - O\_CREAT — create the file if it doesn't exist. This bit has no effect if the file already exists.
  - O\_TRUNC — truncate the file to contain no data if the file exists; this bit has no effect if the file doesn't exist.
  - O\_EXCL — open the file for exclusive access. If the file exists and you also specify O\_CREAT, the open fails (that is, use O\_EXCL to ensure that the file doesn't already exist).

- share* The shared access for the file. This is a combination of the following bits (defined in `<share.h>`):
- SH\_COMPAT — set compatibility mode.
  - SH\_DENYRW — prevent read or write access to the file.
  - SH\_DENYWR — prevent write access to the file.
  - SH\_DENYRD — prevent read access to the file.
  - SH\_DENYNO — permit both read and write access to the file.

If you set O\_CREAT in *oflag*, you must also specify the following argument:

- mode* An object of type `mode_t` that specifies the access mode that you want to use for a newly created file. For more information, see “Access permissions” in the documentation for *stat()*.

## Library:

`libc`

## Description:

The *sopen()* function opens a file at the operating system level for shared access. The name of the file to be opened is given by *filename*.

The file is accessed according to the access mode specified by *oflag*. You must specify O\_CREAT if the file doesn't exist.

The sharing mode of the file is given by the *share* argument. The optional argument is the file permissions to be used when O\_CREAT flag is on in the *oflag* mode; you must provide this when the file is to be created.

The *sopen()* function applies the current file permission mask to the specified permissions (see *umask()*).

Note that

```
open(path, oflag, ...);
```

is the same as:

```
sopen(path, oflag, SH_COMPAT, ...);
```



---

The *sopen()* function ignores advisory locks that you may have set by calling *fcntl()*.

---

## Returns:

A descriptor for the file, or -1 if an error occurs while opening the file (*errno* is set).

## Errors:

|        |                                                                                                                                                                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES | Search permission is denied on a component of the path prefix, or the file exists and the permissions specified by <i>oflag</i> are denied, or the file doesn't exist and write permission is denied for the parent directory of the file to be created. |
| EBUSY  | Sharing mode ( <i>share</i> ) was denied due to a conflicting open.                                                                                                                                                                                      |
| EISDIR | The named file is a directory, and the <i>oflag</i> argument specifies write-only or read/write access.                                                                                                                                                  |
| ELOOP  | Too many levels of symbolic links or prefixes.                                                                                                                                                                                                           |
| EMFILE | No more descriptors available (too many open files).                                                                                                                                                                                                     |
| ENOENT | Path or file not found.                                                                                                                                                                                                                                  |
| ENOSYS | The <i>sopen()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                                                                                                                                              |

**Examples:**

```
#include <stdlib.h>
#include <sys/stat.h>
#include <sys/types.h>
#include <fcntl.h>
#include <share.h>

int main(void)
{
 int filedes ;

 /* open a file for output */
 /* replace existing file if it exists */

 filedes = sopen("file",
 O_WRONLY | O_CREAT | O_TRUNC,
 SH_DENYWR,
 S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP);

 /* read a file which is assumed to exist */

 filedes = sopen("file", O_RDONLY, SH_DENYWR);

 /* append to the end of an existing file */
 /* write a new file if file doesn't exist */

 filedes = sopen("file",
 O_WRONLY | O_CREAT | O_APPEND,
 SH_DENYWR,
 S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP);
 return EXIT_SUCCESS;
}
```

**Classification:**

Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|

**See also:**

*chsize(), close(), creat(), dup(), dup2(), eof(), execl(), execlp(),  
execlpe(), execv(), execve(), execvp(), execvpe(), fcntl(),  
fileno(), fstat(), isatty(), lseek(), open(), read(), stat(), tell(), umask(),  
write()*

## ***sopenfd()***

© 2004, QNX Software Systems Ltd.

*Open for shared access a file associated with a given descriptor*

### **Synopsis:**

```
#include <unistd.h>

int sopenfd(int fd,
 int oflag,
 int sflag);
```

### **Arguments:**

- fd* A file descriptor associated with the file that you want to open.
- oflag* How you want to open the file; a combination of the following bits:
- O\_RDONLY — permit the file to be only read.
  - O\_WRONLY — permit the file to be only written.
  - O\_RDWR — permit the file to be both read and written.
  - O\_APPEND — cause each record that's written to be written at the end of the file.
  - O\_TRUNC — if the file exists, truncate it to contain no data. This flag has no effect if the file doesn't exist.
- sflag* How you want the file to be shared; a combination of the following bits:
- SH\_COMPAT — set compatibility mode.
  - SH\_DENYRW — prevent read or write access to the file.
  - SH\_DENYWR — prevent write access to the file.
  - SH\_DENYRD — prevent read access to the file.
  - SH\_DENYNO — permit both read and write access to the file.



## Library:

`libc`

## Description:

The *sopenfd()* function opens for shared access the file associated with the file descriptor, *fd*. The access mode, *oflag*, must be equal to or more restrictive than the access mode of the source *fd*.

Note that:

```
openfd(fd, oflag);
```

is the same as:

```
sopenfd(fd, oflag, SH_DENYNO);
```

## Returns:

The file descriptor, or -1 if an error occurs (*errno* is set).

## Errors:

- |        |                                                                                                                             |
|--------|-----------------------------------------------------------------------------------------------------------------------------|
| EBADF  | Invalid file descriptor <i>fd</i> .                                                                                         |
| EACCES | The access mode specified by <i>oflag</i> isn't equal to or more restrictive than the access mode of the source <i>fd</i> . |
| EBUSY  | Sharing mode ( <i>sflag</i> ) was denied due to a conflicting open.                                                         |

## Classification:

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*openfd()*

**Synopsis:**

```
#include <spawn.h>

pid_t spawn(const char * path,
 int fd_count,
 const int fd_map[],
 const struct inheritance * inherit,
 char * const argv[],
 char * const envp[]);
```

**Arguments:**

- path*           The full path name of the executable.
- fd\_count*        The number of entries in the *fd\_map* array.
- fd\_map*           An array of file descriptors that you want the child process to inherit. If *fd\_count* isn't 0, *fd\_map* must contain at least *fd\_count* file descriptors, up to OPEN\_MAX FDs. If *fd\_count* is 0, *fd\_map* is ignored.
- When *fdmap[X]* has the value of SPAWN\_FDCLOSED instead of a valid file descriptor, the file descriptor *X* will be closed in the child process.
- If *fd\_count* is 0, all file descriptors (except for the ones created with *fcntl()*'s FD\_CLOEXEC flag) are inherited by the child process.
- inherit*         A structure, of type **struct inheritance**, that indicates what you want the child process to inherit from the parent. This structure contains at least these members:
- unsigned long flags**
- One or more of the following bits:
- SPAWN\_CHECK\_SCRIPT — let *spawn()* start a shell, passing *path* as a script.

- SPAWN\_SEARCH\_PATH — search the **PATH** environment variable for the executable.
- SPAWN\_SETGROUP — set the child's process group to the value in the *pgroup* member. If this flag isn't set, the child process is part of the current process group.
- SPAWN\_SETND — spawn the child process on the node specified by the *nd* member.
- SPAWN\_SETSIGDEF — use the *sigdefault* member to specify the child process's set of defaulted signals. If this flag isn't specified, the child process inherits the parent process's signal actions.
- SPAWN\_SETSIGMASK — use the *sigmask* member to specify the child process's signal mask.

**pid\_t** *pgroup* The child process's group if SPAWN\_SETGROUP is specified in the *flags* member.  
If SPAWN\_SETGROUP is set in *inherit.flags* and *inherit.pgroup* is set to SPAWN\_NEWPGROUP, the child process starts a new process group with the process group ID set to its process ID.

**sigset\_t** *sigmask*  
The child process's signal mask if SPAWN\_SETSIGMASK is specified in the *flags* member.

**sigset\_t sigdefault**

The child process's set of defaulted signals if SPAWN\_SETSIGDEF is specified in the *flags* member.

**uint32\_t nd** The node descriptor of the remote node on which to spawn the child process. This member is used only if SPAWN\_SETND is set in the *flags* member.

*argv* A pointer to an argument vector. The value in *argv[0]* should point to the filename of program being loaded, but can be NULL if no arguments are being passed. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL.

*envp* A pointer to an array of character pointers, each pointing to a string defining an environment variable. The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable. If the value of *envp* is NULL, then the child process inherits the environment of the parent process.

**Library:**

**libc**

**Description:**

The *spawn()* function creates and executes a new child process, named in *path*.



---

If the child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawn()* function is a QNX Neutrino function (based on the POSIX 1003.1d *draft standard*). The C library also includes several specialized *spawn\*()* functions. Their names consist of **spawn** followed by several letters:

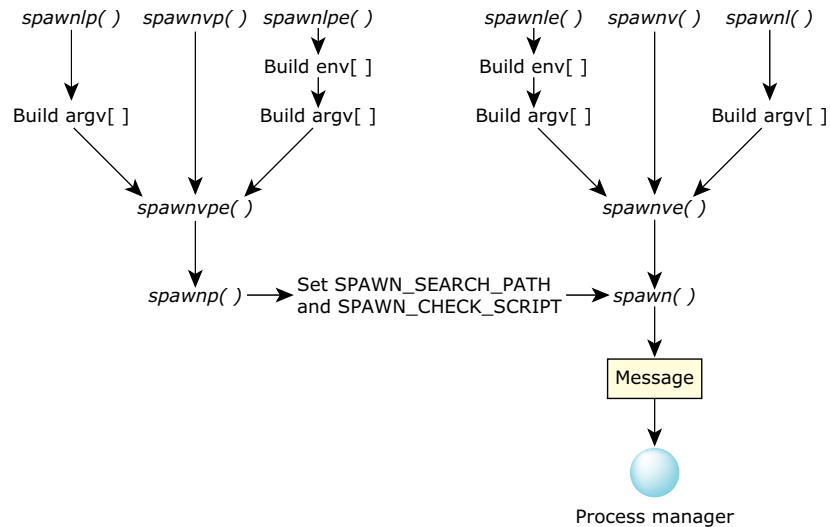
---

| <b>This suffix:</b> | <b>Indicates the function takes these arguments:</b> |
|---------------------|------------------------------------------------------|
|---------------------|------------------------------------------------------|

---

|          |                                                                                                                         |
|----------|-------------------------------------------------------------------------------------------------------------------------|
| <b>e</b> | An array of environment variables.                                                                                      |
| <b>l</b> | A NULL-terminated list of arguments to the program.                                                                     |
| <b>p</b> | A relative path. If the path doesn't contain a slash, the <b>PATH</b> environment variable is searched for the program. |
| <b>v</b> | A vector of arguments to the program.                                                                                   |

As shown below, these functions eventually call *spawn()*, which in turn sends a message to the process manager.



Most of the `spawn*()` functions do a lot of work before a message is sent to `procnto`.

The child process inherits the following attributes of the parent process:

- Process group ID (unless `SPAWN_SETGROUP` is set in *inherit.flags*)
- Session membership
- Real user ID and real group ID
- Supplementary group IDs
- Priority and scheduling policy
- Current working directory and root directory
- File creation mask
- Signal mask (unless `SPAWN_SETSIGMASK` is set in *inherit.flags*)
- Signal actions specified as `SIG_DFL`

- Signal actions specified as SIG\_IGN (except the ones modified by *inherit.sigdefault* when SPAWN\_SETSIGDEF is set in *inherit.flags*)

The child process has several differences from the parent process:

- Signals set to be caught by the parent process are set to the default action (SIG\_DFL).
- The child process's *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* are tracked separately from the parent's.
- The number of seconds left until a SIGALRM signal would be generated is set to zero for the child process.
- The set of pending signals for the child process is empty.
- File locks set by the parent aren't inherited.
- Per-process timers created by the parent aren't inherited.
- Memory locks and mappings set by the parent aren't inherited.

If the child process is spawned on a remote node, the process group ID and the session membership aren't set; the child process is put into a new session and a new process group.

The child process can access the parent process's environment by using the *environ* global variable (found in `<unistd.h>`).

If the *path* is on a filesystem mounted with the ST\_NOSUID flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the child process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the child process is set to the owner ID of *path*. Similarly, if the set-group ID mode bit is set, the effective group ID of the child process is set to the group ID of *path*.

The real user ID, real group ID and supplementary group IDs of the child process remain the same as those of the parent process. The effective user ID and effective group ID of the child process are saved as the saved set-user ID and the saved set-group ID used by the *setuid()*.





---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The process ID of the child process, or -1 if an error occurs (*errno* is set).

## Errors:

|              |                                                                                                                                                                                                                               |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list and environment list of the new child process is greater than ARG_MAX bytes.                                                                                                    |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the child process's file doesn't have the execute bit set or <i>path</i> 's filesystem was mounted with the ST_NOEXEC flag. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                                                                                 |
| EBADF        | An entry in <i>fd_map</i> refers to an invalid file descriptor.                                                                                                                                                               |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                                                                                 |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                                                                |
| EMFILE       | Insufficient resources available to remap file descriptors in the child process.                                                                                                                                              |
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                                                                                                   |
| ENOENT       | The file identified by the <i>path</i> argument is empty, or one or more components of the pathname of the child process don't exist.                                                                                         |

|         |                                                                                                                                                                                                                                  |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENOEXEC | The child process's file has the correct permissions, but isn't in the correct format for an executable. (This error doesn't occur if SPAWN_CHECK_SCRIPT is set in the <i>flags</i> member of the <b>inheritance</b> structure.) |
| ENOMEM  | Insufficient memory available to create the child process.                                                                                                                                                                       |
| ENOSYS  | The <i>spawn()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                                                                                                                      |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                                                                                                                                                           |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *getenv()*, *putenv()*, *setenv()*, *sigaddset()*, *sigdelset()*, *sigemptyset()*, *sigfillset()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvp()*, *spawnvpe()*, *wait()*, *waitpid()*

**Synopsis:**

```
#include <process.h>

int spawnl(int mode,
 const char * path,
 const char * arg0,
 const char * arg1...,
 const char * argn,
 NULL);
```

**Arguments:**

*mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:

- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
- P\_NOWAIT — execute the parent program concurrently with the new child process.
- P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
- P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.

*path* The full path name of the executable.

*arg0, argn, NULL*

The arguments that you want to pass to the new process. You must terminate the list with an argument of NULL.

## Library:

libc

## Description:

The *spawnl()* function creates and executes a new child process, named in *path* with a NULL-terminated list of arguments in *arg0 ... argn*.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnl()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It builds an *argv[ ]* array before calling *spawn()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

Arguments are passed to the child process by supplying one or more pointers to character strings as arguments. These character strings are concatenated with spaces inserted to separate the arguments to form one argument string for the child process. At least one argument, *arg0*, must be passed to the child process. By convention, this first argument is a pointer to the name of the new child process.

The child process inherits the parent's environment. The environment is the collection of environment variables whose values that have been defined with the `export` shell command, the `env` utility, or by the successful execution of the *putenv()* or *setenv()* function. A program may read these values with the *getenv()* function.



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnl()* function's return value depends on the *mode* argument:

| <i>mode</i> | Return value                                                                                                                                               |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|              |                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list of the new child process is greater than ARG_MAX bytes.                                                            |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                   |
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                                      |

|         |                                                                                                                                       |
|---------|---------------------------------------------------------------------------------------------------------------------------------------|
| ENOENT  | The file identified by the <i>path</i> argument is empty, or one or more components of the pathname of the child process don't exist. |
| ENOEXEC | The child process's file has the correct permissions, but isn't in the correct format for an executable.                              |
| ENOMEM  | Insufficient memory available to create the child process.                                                                            |
| ENOSYS  | The <i>spawnl()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                          |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                                                                |

## Examples:

Run `myprog` as if the user had typed:

```
myprog ARG1 ARG2
```

at the command-line:

```
#include <stddef.h>
#include <process.h>

int exit_val;
...
exit_val = spawnl(P_WAIT, "myprog",
 "myprog", "ARG1", "ARG2", NULL);
...
```

The program is found if `myprog` is in the current working directory.

## Classification:

QNX 4

### **Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

### **Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

### **See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(), spawnle(), spawnlp(),  
spawnlpe(), spawnp(), spawnv(), spawnve(), spawnvp(), spawnvpe(),  
wait(), waitpid()*

## ***spawnle()***

© 2004, QNX Software Systems Ltd.

*Spawn a child process, given a list of arguments and an environment*

### **Synopsis:**

```
#include <process.h>

int spawnle(int mode,
 const char * path,
 const char * arg0,
 const char * arg1...,
 const char * argn,
 NULL,
 const char * envp[]);
```

### **Arguments:**

- mode*     How you want to load the child process, and how you want the parent program to behave after the child program is initiated:
- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
  - P\_NOWAIT — execute the parent program concurrently with the new child process.
  - P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
  - P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.
- path*     The full path name of the executable.
- arg0, argn, NULL*
- The arguments that you want to pass to the new process. You must terminate the list with an argument of NULL.
- envp*     NULL, or a pointer to an array of character pointers, each pointing to a string that defines an environment variable.



The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable.

## Library:

`libc`

## Description:

The *spawnle()* function creates and executes a new child process, named in *path* with NULL-terminated list of arguments in *arg0* ... *argn* and with the environment specified in *envp*.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnle()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It builds *argv[ ]* and *envp[ ]* arrays before calling *spawn()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

Arguments are passed to the child process by supplying one or more pointers to character strings as arguments. These character strings are concatenated with spaces inserted to separate the arguments to form one argument string for the child process. At least one argument, *arg0*, must be passed to the child process. By convention, this first argument is a pointer to the name of the new child process.

If *envp* is NULL, the child process inherits the environment of the parent process. The new process can access the calling process environment by using the *environ* global variable (found in `<unistd.h>`).



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnle()* function's return value depends on the *mode* argument:

| <i>mode</i> | <b>Return value</b>                                                                                                                                        |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|         |                                                                                                                                                                  |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG   | The number of bytes used by the argument list or environment list of the new child process is greater than ARG_MAX bytes.                                        |
| EACCESS | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN  | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT  | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP   | Too many levels of symbolic links or prefixes.                                                                                                                   |

|              |                                                                                                                                       |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                           |
| ENOENT       | The file identified by the <i>path</i> argument is empty, or one or more components of the pathname of the child process don't exist. |
| ENOEXEC      | The child process's file has the correct permissions, but isn't in the correct format for an executable.                              |
| ENOMEM       | Insufficient memory available to create the child process.                                                                            |
| ENOSYS       | The <i>spawnle()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                         |
| ENOTDIR      | A component of the path prefix of the child process isn't a directory.                                                                |

## Examples:

Run **myprog** as if the user had typed:

```
myprog ARG1 ARG2
```

at the command-line:

```
#include <stddef.h>
#include <process.h>

char *env_list[] = { "SOURCE=MYDATA",
 "TARGET=OUTPUT",
 "lines=65",
 NULL
 };

spawnle(P_WAIT, "myprog",
 "myprog", "ARG1", "ARG2", NULL,
 env_list);
```

The program is found if **myprog** is in the current working directory. The environment for the child program consists of the three environment variables **SOURCE**, **TARGET** and **lines**.

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl(), execlp(), execlpe(), execv(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(), spawnl(), spawnlp(),  
spawnlpe(), spawnp(), spawnv(), spawnve(), spawnvp(), spawnvpe(),  
wait(), waitpid()*

*Spawn a child process, given a list of arguments and a relative path*

**Synopsis:**

```
#include <process.h>

int spawnlp(int mode,
 const char * file,
 const char * arg0,
 const char * arg1...,
 const char * argn,
 NULL);
```

**Arguments:**

*mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:

- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
- P\_NOWAIT — execute the parent program concurrently with the new child process.
- P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
- P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.

*file* The name of the executable file. If this argument contains a slash, it's used as the pathname of the executable; otherwise, the function searches for *file* in the directories listed in the **PATH** environment variable.

*arg0, argn, NULL*

The arguments that you want to pass to the new process. You must terminate the list with an argument of NULL.

## Library:

libc

## Description:

The *spawnlp()* function creates and executes a new child process, named in *file* with NULL-terminated list of arguments in *arg0 ... argn*.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnlp()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It builds an *argv[ ]* array before calling *spawnp()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

Arguments are passed to the child process by supplying one or more pointers to character strings as arguments. These character strings are concatenated with spaces inserted to separate the arguments to form one argument string for the child process. At least one argument, *arg0*, must be passed to the child process. By convention, this first argument is a pointer to the name of the new child process.

The child process inherits the parent's environment. The environment is the collection of environment variables whose values that have been defined with the `export` shell command, the `env` utility, or by the successful execution of the *putenv()* or *setenv()* function. A program may read these values with the *getenv()* function.



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnlp()* function's return value depends on the *mode* argument:

| <i>mode</i> | Return value                                                                                                                                               |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|              |                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list of the new child process is greater than ARG_MAX bytes.                                                            |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                   |
| ENAMETOOLONG | The length of <i>file</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                                      |

|         |                                                                                                                                       |
|---------|---------------------------------------------------------------------------------------------------------------------------------------|
| ENOENT  | The file identified by the <i>file</i> argument is empty, or one or more components of the pathname of the child process don't exist. |
| ENOEXEC | The child process's file has the correct permissions, but isn't in the correct format for an executable.                              |
| ENOMEM  | Insufficient memory available to create the child process.                                                                            |
| ENOSYS  | The <i>spawnlp()</i> function isn't implemented for the filesystem specified in <i>file</i> .                                         |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                                                                |

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(), spawnl(), spawnle(),  
spawnlpe(), spawnlp(), spawnp(), spawnv(), spawnve(), spawnvp(), spawnvpe(),  
wait(), waitpid()*



*Spawn a child process, given a list of arguments, an environment, and a relative path*

**Synopsis:**

```
#include <process.h>

int spawnlpe(int mode,
 const char * file,
 const char * arg0,
 const char * arg1...,
 const char * argn,
 NULL,
 const char * envp[]);
```

**Arguments:**

*mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:

- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
- P\_NOWAIT — execute the parent program concurrently with the new child process.
- P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
- P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.

*file* The name of the executable file. If this argument contains a slash, it's used as the pathname of the executable; otherwise, the function searches for *file* in the directories listed in the **PATH** environment variable.

*arg0, argn, NULL*

The arguments that you want to pass to the new process. You must terminate the list with an argument of NULL.

*envp* NULL, or a pointer to an array of character pointers, each pointing to a string that defines an environment variable. The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable.

## Library:

`libc`

## Description:

The *spawnlpe()* function creates and executes a new child process, named in *file* with NULL-terminated list of arguments in *arg0 ... argn* and with the environment specified in *envp*.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnlpe()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It builds *argv[ ]* and *envp[ ]* arrays before calling *spawnp()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

Arguments are passed to the child process by supplying one or more pointers to character strings as arguments. These character strings are concatenated with spaces inserted to separate the arguments to form one argument string for the child process. At least one argument, *arg0*, must be passed to the child process. By convention, this first argument is a pointer to the name of the new child process.

If the value of *envp* is NULL, then the child process inherits the environment of the parent process. The new process can access the

calling process environment by using the *environ* global variable (found in `<unistd.h>`).



A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnlpe()* function's return value depends on the *mode* argument:

| <i>mode</i> | Return value                                                                                                                                               |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|         |                                                                                                                                                                  |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG   | The number of bytes used by the argument list or environment list of the new child process is greater than ARG_MAX bytes.                                        |
| EACCESS | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN  | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT  | One of the buffers specified in the function call is invalid.                                                                                                    |

|              |                                                                                                                                       |
|--------------|---------------------------------------------------------------------------------------------------------------------------------------|
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                        |
| ENAMETOOLONG | The length of <i>file</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                           |
| ENOENT       | The file identified by the <i>file</i> argument is empty, or one or more components of the pathname of the child process don't exist. |
| ENOEXEC      | The child process's file has the correct permissions, but isn't in the correct format for an executable.                              |
| ENOMEM       | Insufficient memory available to create the child process.                                                                            |
| ENOSYS       | The <i>spawnlpe()</i> function isn't implemented for the filesystem specified in <i>file</i> .                                        |
| ENOTDIR      | A component of the path prefix of the child process isn't a directory.                                                                |

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl(), execlp(), execlpe(), execv(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(),  
spawnl(), spawnle(), spawnlp(), spawnnp(),  
spawnv(), spawnve(), spawnvp(), spawnvpe(),  
wait(), waitpid()*

## ***spawnp()***

© 2004, QNX Software Systems Ltd.

*Create and execute a new child process, given a relative path*

### **Synopsis:**

```
#include <spawn.h>

pid_t spawnp(const char * file,
 int fd_count,
 const int fd_map[],
 const struct inheritance * inherit,
 char * const argv[],
 char * const envp[]);
```

### **Arguments:**

- file* If this argument contains a slash, it's used as the pathname of the executable; otherwise, the **PATH** environment variable is searched for *file*.
- fd\_count* The number of entries in the *fd\_map* array.
- fd\_map* An array of file descriptors that you want the child process to inherit. If *fd\_count* isn't 0, *fd\_map* must contain at least *fd\_count* file descriptors, up to OPEN\_MAX FDs. If *fd\_count* is 0, *fd\_map* is ignored.
- When *fdmap[X]* has the value of SPAWN\_FDCLOSED instead of a valid file descriptor, the file descriptor *X* will be closed in the child process.
- If *fd\_count* is 0, all file descriptors (except for the ones created with *fcntl()*'s FD\_CLOEXEC flag) are inherited by the child process.
- inherit* A structure, of type **struct inheritance**, that indicates what you want the child process to inherit from the parent. This structure contains at least these members:

**unsigned long** *flags*

One or more of the following bits:

- SPAWN\_CHECK\_SCRIPT — let *spawn()* start a shell, passing *path* as a script.
- SPAWN\_SEARCH\_PATH — search the **PATH** environment variable for the executable.
- SPAWN\_SETGROUP — set the child process's group to the value in the *pgroup* member. If this flag isn't set, the child process is part of the current process group.
- SPAWN\_SETND — spawn the child process on the node specified by the *nd* member.
- SPAWN\_SETSIGDEF — use the *sigdefault* member to specify the child process's set of defaulted signals. If this flag isn't specified, the child process inherits the parent process's signal actions.
- SPAWN\_SETSIGMASK — use the *sigmask* member to specify the child process's signal mask.

**pid\_t** *pgroup* The child process's group if SPAWN\_SETGROUP is specified in the *flags* member.  
If SPAWN\_SETGROUP is set in *inherit.flags* and *inherit.pgroup* is set to SPAWN\_NEWPGROUP, the child process starts a new process group with the process group ID set to its process ID.

**sigset\_t** *sigmask*  
The child process's signal mask if SPAWN\_SETSIGMASK is specified in the *flags* member.

**sigset\_t sigdefault**

The child process's set of defaulted signals if SPAWN\_SETSIGDEF is specified in the *flags* member.

**uint32\_t nd** The node descriptor of the remote node on which to spawn the child process. This member is used only if SPAWN\_SETND is set in the *flags* member.

**argv** A pointer to an argument vector. The value in *argv[0]* should point to the filename of program being loaded, but can be NULL if no arguments are being passed. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL.

**envp** A pointer to an array of character pointers, each pointing to a string defining an environment variable. The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable. If the value of *envp* is NULL, then the child process inherits the environment of the parent process.

## Library:

libc

## Description:

The *spawnp()* function creates and executes a new child process, named in *file*. It sets the SPAWN\_CHECK\_SCRIPT and SPAWN\_SEARCH\_PATH flags (see below), and calls *spawn()*.






---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnp()* function is a QNX function (based on the POSIX 1003.1d *draft standard*). The C library also includes several specialized *spawn\*()* functions. Their names consist of **spawn** followed by several letters:

---

**This suffix:    Indicates the function takes these arguments:**

---

|          |                                                                                                                         |
|----------|-------------------------------------------------------------------------------------------------------------------------|
| <b>e</b> | An array of environment variables.                                                                                      |
| <b>l</b> | A NULL-terminated list of arguments to the program.                                                                     |
| <b>p</b> | A relative path. If the path doesn't contain a slash, the <b>PATH</b> environment variable is searched for the program. |
| <b>v</b> | A vector of arguments to the program.                                                                                   |

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

The child process inherits the following attributes of the parent process:

- Process group ID (unless `SPAWN_SETGROUP` is set in *inherit.flags*)
- Session membership
- Real user ID and real group ID
- Supplementary group IDs
- Priority and scheduling policy
- Current working directory and root directory

- File creation mask
- Signal mask (unless SPAWN\_SETSIGMASK is set in *inherit.flags*)
- Signal actions specified as SIG\_DFL
- Signal actions specified as SIG\_IGN (except the ones modified by *inherit.sigdefault* when SPAWN\_SETSIGDEF is set in *inherit.flags*)

The child process has several differences from the parent process:

- Signals set to be caught by the parent process are set to the default action (SIG\_DFL).
- The child process's *tms\_utime*, *tms\_stime*, *tms\_cutime*, and *tms\_cstime* are tracked separately from the parent's.
- The number of seconds left until a SIGALRM signal would be generated is set to zero for the child process.
- The set of pending signals for the child process is empty.
- File locks set by the parent aren't inherited.
- Per-process timers created by the parent aren't inherited.
- Memory locks and mappings set by the parent aren't inherited.

If the child process is spawned on a remote node, the process group ID and the session membership aren't set; the child process is put into a new session and a new process group.

The child process can access the parent process's environment by using the *environ* global variable (found in `<unistd.h>`).

If the *file* is on a filesystem mounted with the ST\_NOSUID flag set, the effective user ID, effective group ID, saved set-user ID and saved set-group ID are unchanged for the child process. Otherwise, if the set-user ID mode bit is set, the effective user ID of the child process is set to the owner ID of *file*. Similarly, if the set-group ID mode bit is set, the effective group ID of the child process is set to the group ID of *file*. The real user ID, real group ID and supplementary group IDs

of the child process remain the same as those of the parent process. The effective user ID and effective group ID of the child process are saved as the saved set-user ID and the saved set-group ID used by the *setuid()*.



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The process ID of the child process, or -1 if an error occurs (*errno* is set).

## Errors:

|         |                                                                                                                                                                                                                                   |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG   | The number of bytes used by the argument list and environment list of the new child process is greater than ARG_MAX bytes.                                                                                                        |
| EACCESS | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set or <i>file</i> 's filesystem was mounted with the ST_NOEXEC flag. |
| EAGAIN  | Insufficient resources available to create the child process.                                                                                                                                                                     |
| EBADF   | An entry in <i>fd_map</i> refers to an invalid file descriptor.                                                                                                                                                                   |
| EFAULT  | One of the buffers specified in the function call is invalid.                                                                                                                                                                     |
| ELOOP   | Too many levels of symbolic links or prefixes.                                                                                                                                                                                    |
| EMFILE  | Insufficient resources available to remap file descriptors in the child process.                                                                                                                                                  |

ENAMETOOLONG

The length of *file* exceeds PATH\_MAX or a pathname component is longer than NAME\_MAX.

ENOENT

The file identified by the *file* argument is empty, or one or more components of the pathname of the child process don't exist.

ENOEXEC

The child process's file has the correct permissions, but isn't in the correct format for an executable.

ENOMEM

Insufficient memory available to create the child process.

ENOSYS

The *spawnp()* function isn't implemented for the filesystem specified in *file*.

ENOTDIR

A component of the path prefix of the child process isn't a directory.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), sigaddset(), sigdelset(),  
sigemptyset(), sigfillset(), spawn(), spawnl(), spawnle(), spawnlp(),*

*spawnlpe(), spawnv(), spawnve(), spawnvp(), spawnvpe(), wait(),  
waitpid()*

## ***spawnv()***

© 2004, QNX Software Systems Ltd.

*Spawn a child process, given a vector of arguments*

### **Synopsis:**

```
#include <process.h>

int spawnv(int mode,
 const char * path,
 char * const argv[]);
```

### **Arguments:**

- mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:
- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
  - P\_NOWAIT — execute the parent program concurrently with the new child process.
  - P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
  - P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.
- path* The full path name of the executable.
- argv* A pointer to an argument vector. The value in *argv[0]* should point to a filename that's associated with the program that you're loading. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL, and *argv[0]* can't be a NULL pointer, even if you're not passing any argument strings.

## Library:

`libc`

## Description:

The *spawnv()* function creates and executes a new child process, named in *path* with the NULL-terminated list of arguments in the *argv* vector.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnv()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It calls *spawnve()* before calling *spawn()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

The child process inherits the parent's environment. The environment is the collection of environment variables whose values that have been defined with the `export` shell command, the `env` utility, or by the successful execution of the *putenv()* or *setenv()* function. A program may read these values with the *getenv()* function.



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnv()* function's return value depends on the *mode* argument:

| <i>mode</i> | <b>Return value</b>                                                                                                                                        |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

### Errors:

|              |                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list of the new child process is greater than ARG_MAX bytes.                                                            |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                   |
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                                      |
| ENOENT       | The file identified by the <i>path</i> argument is empty, or one or more components of the pathname of the child process don't exist.                            |
| ENOEXEC      | The child process's file has the correct permissions, but isn't in the correct format for an executable.                                                         |



|         |                                                                                              |
|---------|----------------------------------------------------------------------------------------------|
| ENOMEM  | Insufficient memory available to create the child process.                                   |
| ENOSYS  | The <i>spawnv()</i> function isn't implemented for the filesystem specified in <i>path</i> . |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                       |

**Examples:**

Run **myprog** as if a user had typed:

```
myprog ARG1 ARG2
```

at the command-line:

```
#include <stddef.h>
#include <process.h>

char *arg_list[] = { "myprog", "ARG1", "ARG2", NULL };
:
spawnv(P_WAIT, "myprog", arg_list);
```

The program is found if **myprog** is in the current working directory.

**Classification:**

QNX 4

**Safety**


---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(), spawnl(), spawnle(),  
spawnlp(), spawnlpe(), spawnp(), spawnve(), spawnvp(), spawnvpe(),  
wait(), waitpid()*

*Spawn a child process, given a vector of arguments and an environment*

## Synopsis:

```
#include <process.h>

int spawnve(int mode,
 const char * path,
 char * const argv[],
 char * const envp[]);
```

## Arguments:

- mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:
- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
  - P\_NOWAIT — execute the parent program concurrently with the new child process.
  - P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
  - P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.
- path* The full path name of the executable.
- argv* A pointer to an argument vector. The value in *argv[0]* should point to a filename that's associated with the program that you're loading. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL, and *argv[0]* can't be a NULL pointer, even if you're not passing any argument strings.
- envp* NULL, or a pointer to an array of character pointers, each pointing to a string that defines an environment variable.

The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable.

## Library:

`libc`

## Description:

The *spawnve()* function creates and executes a new child process, named in *path* with the NULL-terminated list of arguments in the *argv* vector.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnve()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It calls *spawn()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

If the value of *envp* is NULL, then the child process inherits the environment of the parent process. The new process can access the calling process environment by using the *environ* global variable (found in `<unistd.h>`).



---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnve()* function's return value depends on the *mode* argument:

| <i>mode</i> | <b>Return value</b>                                                                                                                                        |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You can't get the exit status of a P_NOWAITO process.          |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|              |                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list or environment list of the new child process is greater than ARG_MAX bytes.                                        |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                   |
| ENAMETOOLONG | The length of <i>path</i> exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                                      |

|         |                                                                                                                                       |
|---------|---------------------------------------------------------------------------------------------------------------------------------------|
| ENOENT  | The file identified by the <i>path</i> argument is empty, or one or more components of the pathname of the child process don't exist. |
| ENOEXEC | The child process's file has the correct permissions, but isn't in the correct format for an executable.                              |
| ENOMEM  | Insufficient memory available to create the child process.                                                                            |
| ENOSYS  | The <i>spawnve()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                         |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                                                                |

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *getenv()*, *putenv()*, *setenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnvp()*, *spawnvpe()*, *wait()*, *waitpid()*

*Spawn a child process, given a vector of arguments and a relative path*

## Synopsis:

```
#include <process.h>

int spawnvp(int mode,
 const char * file,
 char * const argv[]);
```

## Arguments:

- mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:
- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
  - P\_NOWAIT — execute the parent program concurrently with the new child process.
  - P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
  - P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.
- file* The name of the executable file. If this argument contains a slash, it's used as the pathname of the executable; otherwise, the function searches for *file* in the directories listed in the **PATH** environment variable.
- argv* A pointer to an argument vector. The value in *argv[0]* should point to a filename that's associated with the program that you're loading. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL, and *argv[0]* can't be a NULL pointer, even if you're not passing any argument strings.

## Library:

`libc`

## Description:

The *spawnvp()* function creates and executes a new child process, named in *file* with the NULL-terminated list of arguments in the *argv* vector.



If the new process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnvp()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It calls *spawnvpe()* before calling *spawnvp()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

The child process inherits the parent's environment. The environment is the collection of environment variables whose values that have been defined with the `export` shell command, the `env` utility, or by the successful execution of the *putenv()* or *setenv()* function. A program may read these values with the *getenv()* function.



A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnvp()* function's return value depends on the *mode* argument:



| <i>mode</i> | <b>Return value</b>                                                                                                                                        |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You cannot get the exit status of a P_NOWAITO process.         |

If an error occurs, -1 is returned (*errno* is set).

## **Errors:**

|              |                                                                                                                                                                  |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG        | The number of bytes used by the argument list of the new child process is greater than ARG_MAX bytes.                                                            |
| EACCESS      | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process's file doesn't have the execute bit set. |
| EAGAIN       | Insufficient resources available to create the child process.                                                                                                    |
| EFAULT       | One of the buffers specified in the function call is invalid.                                                                                                    |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                   |
| ENAMETOOLONG | The length of <i>file</i> and its path exceeds PATH_MAX or a pathname component is longer than NAME_MAX.                                                         |
| ENOENT       | The file identified by the <i>file</i> argument is empty, or one or more components of the pathname of the new process don't exist.                              |
| ENOEXEC      | The child process file has the correct permissions, but isn't in the correct format for an executable.                                                           |

|         |                                                                                               |
|---------|-----------------------------------------------------------------------------------------------|
| ENOMEM  | Insufficient memory available to create the child process.                                    |
| ENOSYS  | The <i>spawnvp()</i> function isn't implemented for the filesystem specified in <i>file</i> . |
| ENOTDIR | A component of the path prefix of the child process isn't a directory.                        |

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl()*, *execle()*, *execlp()*, *execlpe()*, *execv()*, *execve()*, *execvp()*, *execvpe()*, *getenv()*, *putenv()*, *setenv()*, *spawn()*, *spawnl()*, *spawnle()*, *spawnlp()*, *spawnlpe()*, *spawnp()*, *spawnv()*, *spawnve()*, *spawnvpe()*, *wait()*, *waitpid()*

Spawn a child process, given a vector of arguments, an environment, and a relative path

## Synopsis:

```
#include <spawn.h>

int spawnvpe(int mode,
 const char * file,
 char * const argv[],
 char * const envp[]);
```

## Arguments:

- mode* How you want to load the child process, and how you want the parent program to behave after the child program is initiated:
- P\_WAIT — load the child program into available memory, execute it, and make the parent program resume execution after the child process ends.
  - P\_NOWAIT — execute the parent program concurrently with the new child process.
  - P\_NOWAITO — execute the parent program concurrently with the new child process. You can't use *wait()* to obtain the exit code.
  - P\_OVERLAY — replace the parent program with the child program in memory and execute the child. No return is made to the parent program. This is equivalent to calling the appropriate *exec\*()* function.
- file* The name of the executable file. If this argument contains a slash, it's used as the pathname of the executable; otherwise, the function searches for *file* in the directories listed in the **PATH** environment variable.
- argv* A pointer to an argument vector. The value in *argv[0]* should point to a filename that's associated with the program that you're loading. The last member of *argv* must be a NULL pointer. The value of *argv* can't be NULL, and *argv[0]* can't be a NULL pointer, even if you're not passing any argument strings.

*envp* NULL, or a pointer to an array of character pointers, each pointing to a string that defines an environment variable. The array is terminated with a NULL pointer. Each pointer points to a character string of the form:

*variable=value*

that's used to define an environment variable.

## Library:

`libc`

## Description:

The *spawnvpe()* function creates and executes a new child process, named in *file* with the NULL-terminated list of arguments in the *argv* vector.



---

If the new child process is a shell script, the first line must start with `#!`, followed by the path and arguments of the shell to be run to interpret the script. The script must also be marked as executable.

---

The *spawnvpe()* function *isn't* a POSIX 1003.1 function, and isn't guaranteed to behave the same on all operating systems. It calls *spawnp()*.

For a diagram of how the *spawn\** functions are related, see the description of *spawn()*.

If the value of *envp* is NULL, then the child process inherits the environment of the parent process. The new process can access the calling process environment by using the *environ* global variable (found in `<unistd.h>`).




---

A parent/child relationship doesn't imply that the child process dies when the parent process dies.

---

## Returns:

The *spawnvpe()* function's return value depends on the *mode* argument:

| <i>mode</i> | Return value                                                                                                                                               |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| P_WAIT      | The exit status of the child process.                                                                                                                      |
| P_NOWAIT    | The process ID of the child process. To get the exit status for a P_NOWAIT process, you must use the <i>waitpid()</i> function, giving it this process ID. |
| P_NOWAITO   | The process ID of the child process, or 0 if the process is being started on a remote node. You cannot get the exit status of a P_NOWAITO process.         |

If an error occurs, -1 is returned (*errno* is set).

## Errors:

|         |                                                                                                                                                                |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| E2BIG   | The number of bytes used by the argument list or environment list of the new child process is greater than ARG_MAX bytes.                                      |
| EACCESS | Search permission is denied for a directory listed in the path prefix of the new child process or the new child process file doesn't have the execute bit set. |
| EAGAIN  | Insufficient resources available to create the child process.                                                                                                  |
| EFAULT  | One of the buffers specified in the function call is invalid.                                                                                                  |
| ELOOP   | Too many levels of symbolic links or prefixes.                                                                                                                 |

ENAMETOOLONG

The length of *file* plus its path exceeds PATH\_MAX or a pathname component is longer than NAME\_MAX.

ENOENT

The file identified by the *file* argument is empty, or one or more components of the pathname of the child process don't exist.

ENOEXEC

The child process file has the correct permissions, but isn't in the correct format for an executable.

ENOMEM

Insufficient memory available to create the child process.

ENOSYS

The *spawnvpe()* function isn't implemented for the filesystem specified in *file*.

ENOTDIR

A component of the path prefix of the child process isn't a directory.

**Classification:**

QNX 4

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Read the <i>Caveats</i> |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Yes                     |

**Caveats:**

If *mode* is P\_WAIT, this function is a cancellation point.

**See also:**

*execl(), execlp(), execlpe(), execv(), execve(), execvp(),  
execvpe(), getenv(), putenv(), setenv(), spawn(), spawnl(), spawnle(),  
spawnlp(), spawnlpe(), spawnp(), spawnv(), spawnve(), spawnvp(),  
wait(), waitpid()*

# ***sprintf()***

© 2004, QNX Software Systems Ltd.

*Print formatted output into a string*

## **Synopsis:**

```
#include <stdio.h>

int sprintf(char* buf,
 const char* format,
 ...);
```

## **Arguments:**

*buf*            A pointer to the buffer where you want to function to store the formatted string.

*format*        A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

## **Library:**

libc

## **Description:**

The *sprintf()* function is similar to *fprintf()*, except that *sprintf()* places the generated output into the character array pointed to by *buf*, instead of writing it to a file. A null character is placed at the end of the generated character string.

## **Returns:**

The number of characters written into the array, not counting the terminating null character. An error can occur while converting a value for output. When an error occurs, *errno* indicates the type of error detected.

## **Examples:**

```
#include <stdio.h>
#include <stdlib.h>

/* Create temporary file names using a counter */
```



```
char namebuf[13];
int TempCount = 0;

char* make_temp_name()
{
 sprintf(namebuf, "ZZ%.6o.TMP", TempCount++);
 return(namebuf);
}

int main(void)
{
 FILE* tf1,* tf2;

 tf1 = fopen(make_temp_name(), "w");
 tf2 = fopen(make_temp_name(), "w");
 fputs("temp file 1", tf1);
 fputs("temp file 2", tf2);
 fclose(tf1);
 fclose(tf2);
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *swprintf()*, *vfprintf()*, *vwprintf()*, *vprintf()*, *vsprintf()*, *vsnprintf()*, *vswprintf()*, *wprintf()*

## ***sqrt(), sqrtf()***

© 2004, QNX Software Systems Ltd.

*Calculate the nonnegative square root of a number*

### **Synopsis:**

```
#include <math.h>

double sqrt(double x);

float sqrtf(float x);
```

### **Arguments:**

*x*     The number that you want to calculate the square root of.

### **Library:**

`libm`

### **Description:**

These functions compute the nonnegative square root of *x*. A domain error occurs if the argument is negative.

### **Returns:**

The nonnegative square root of the given number.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

### **Examples:**

```
#include <stdio.h>
#include <math.h>
#include <stdlib.h>

int main(void)
{
 printf("%f\n", sqrt(.5));
 return EXIT_SUCCESS;
}
```

produces the output:

0.707107

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

*errno, exp(), log(), pow()*

## ***srand()***

© 2004, QNX Software Systems Ltd.

*Start a new sequence of pseudo-random integers*

### **Synopsis:**

```
#include <stdlib.h>

void srand(unsigned int seed);
```

### **Arguments:**

*seed*     The seed of the sequence of pseudo-random integers.

### **Library:**

libc

### **Description:**

The *srand()* function uses the argument *seed* to start a new sequence of pseudo-random integers to be returned by subsequent calls to *rand()*. A particular sequence of pseudo-random integers can be repeated by calling *srand()* with the same *seed* value. The default sequence of pseudo-random integers is selected with a *seed* value of 1.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

int main(void)
{
 int i;

 srand(982);
 for(i = 1; i < 10; ++i) {
 printf("%d\n", rand());
 }

 /* Start the same sequence over again. */

 srand(982);
 for(i = 1; i < 10; ++i) {
 printf("%d\n", rand());
 }
}
```

```
/*
 Use the current time as a seed to
 get a different sequence.
*/

srand((int) time(NULL));
for(i = 1; i < 10; ++i) {
 printf("%d\n", rand());
}

return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*rand()*

## ***srand48()***

© 2004, QNX Software Systems Ltd.

*Initialize a sequence of pseudo-random numbers*

### **Synopsis:**

```
#include <stdlib.h>

void srand48(long seed);
```

### **Arguments:**

*seed*     The seed of the sequence of pseudo-random integers.

### **Library:**

`libc`

### **Description:**

The *srand48()* is used to initialize the internal buffer  $r(n)$  of *drand48()*, *lrand48()*, and *mrnd48()* such that the 32 bits of the seed value are copied into the upper 32 bits of  $r(n)$ , with the lower 16 bits of  $r(n)$  arbitrarily being set to `0x330E`. Additionally, the constant multiplicand and addend of the algorithm are reset to the default values: the multiplicand  $a = 0xFDEECE66D = 25214903917$  and the addend  $c = 0xB = 11$ .

### **Classification:**

Standard Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*drand48()*, *erand48()*, *jrand48()*, *lcong48()*, *lrand48()*, *mrnd48()*,  
*nrand48()*, *seed48()*

## ***srandom()***

© 2004, QNX Software Systems Ltd.

*Set the seed for a pseudo-random number generator*

### **Synopsis:**

```
#include <stdlib.h>

void srandom(unsigned int seed);
```

### **Arguments:**

*seed*     The seed of the sequence of pseudo-random integers.

### **Library:**

`libc`

### **Description:**

The *srandom()* function initializes the current state array using the value of *seed*.

Use this function in conjunction with the following:

|                    |                                                             |
|--------------------|-------------------------------------------------------------|
| <i>initstate()</i> | Initialize the state of the pseudo-random number generator. |
| <i>random()</i>    | Generate a pseudo-random number using a default state.      |
| <i>setstate()</i>  | Specify the state of the pseudo-random number generator.    |

The *random()* and *srandom()* functions have (almost) the same calling sequence and initialization properties as *rand()* and *srand()*. Unlike *srand()*, *srandom()* doesn't return the old seed because the amount of state information used is much more than a single word. The *initstate()* and *setstate()* routines are provided to deal with restarting/changing random number generators. With 256 bytes of state information, the period of the random-number generator is greater than 269.



Like *rand()*, *random()* produces by default a sequence of numbers that can be duplicated by calling *srandom()* with 1 as the seed.

After initialization, a state array can be restarted at a different point in one of two ways:

- The *initstate()* function can be used, with the desired seed, state array, and size of the array.
- The *setstate()* function, with the desired state, can be used, followed by *srandom()* with the desired seed. The advantage of using both of these functions is that the size of the state array does not have to be saved once it is initialized.

## Classification:

Standard Unix

### Safety

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*drand48()*, *initstate()*, *rand()*, *random()*, *setstate()*, *srand()*

## **`_srealloc()`**

© 2004, QNX Software Systems Ltd.

*Allocate, reallocate or free a block of memory*

### **Synopsis:**

```
#include <malloc.h>

void *_srealloc(void* ptr,
 size_t old_size,
 size_t new_size);
```

### **Arguments:**

*ptr*            NULL, or a pointer to the block of memory that you want to reallocate.

*old\_size*      The current size of the block, in bytes.

*new\_size*      The size of the block to allocate, in bytes.

### **Library:**

`libc`

### **Description:**

When the value of the *ptr* argument is NULL, a new block of memory of *new\_size* bytes is allocated.

If the value of *new\_size* is zero, the corresponding `_sfree()` function is called to release *old\_size* bytes of memory memory pointed to by *ptr*.

Otherwise, the `_srealloc()` function reallocates space for an object of *new\_size* bytes by doing one of the following:

- Shrinking the allocated size of the allocated memory block *ptr* when *new\_size* is sufficiently smaller than *old\_size*.
- Or:
- Extending the allocated size of the allocated memory block *ptr* if there is a large enough block of unallocated memory immediately following *ptr*.
- Or:

- Allocating a new block, and copying the contents of *ptr* to the new block.




---

Because it's possible that a new block will be allocated, don't maintain any pointers into the old memory after a successful call to this function. These pointers will point to freed memory, with possible disastrous results when a new block is allocated.

---

The function returns `NULL` when the memory pointed to by *ptr* can't be reallocated. In this case, the memory pointed to by *ptr* isn't freed, so be sure to keep a pointer to the old memory block.

```
buffer = (char *) _srealloc(buffer, 100, 200);
```

In the above example, *buffer* is set to `NULL` if the function fails, and no longer points to the old memory block. If *buffer* is your only pointer to the memory block, then you've lost access to this memory.

The *\_srealloc()* function reallocates memory from the heap.

You must use *\_sfree()* to deallocate the memory allocated by *\_srealloc()*.

## Returns:

A pointer to the start of the reallocated memory, or `NULL` if there's insufficient memory available, or if the value of the *new\_size* argument is zero.

## Classification:

QNX Neutrino

### Safety

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*calloc()*, *free()*, *realloc()*, *\_scalloc()*, *\_sfree()*, *\_smalloc()*

## Synopsis:

```
#include <stdio.h>

int sscanf(const char* in_string,
 const char* format,
 ...);
```

## Arguments:

*in\_string*     The string that you want to read from.

*format*        A string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

## Library:

libc

## Description:

The *sscanf()* function scans input from the character string *in\_string*, under control of the argument *format*. Following the format string is the list of addresses of items to receive values.

## Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF when the scanning is terminated by reaching the end of the input string.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int day, year;
 char weekday[20], month[20];
```

```
 sscanf("Thursday February 0025 1999",
 "%s %s %d %d",
 weekday, month, &day, &year);
 printf("%s %d, %d is a %s\n",
 month, day, year, weekday);
 return EXIT_SUCCESS;
}
```

produces the following:

```
February 25, 1999 is a Thursday
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*fscanf()*, *fwscanf()*, *scanf()*, *swscanf()*, *vfscanf()*, *vfwscanf()*, *vscanf()*,  
*vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

## Synopsis:

```
#include <sys/stat.h>

int stat(const char * path,
 struct stat * buf);

int stat64(const char * path,
 struct stat64 * buf);
```

## Arguments:

*path*     The path of the file or directory that you want information about.

*buf*      A pointer to a buffer where the function can store the information; see below.

## Library:

libc

## Description:

The *stat()* and *stat64()* functions obtain information about the file or directory referenced in *path*. This information is placed in the structure located at the address indicated by *buf*.

### **stat structure**

Here's the **stat** structure that's defined in **<sys/stat.h>**:

```

struct stat {
#if _FILE_OFFSET_BITS - 0 == 64
 ino_t st_ino; /* File serial number. */
 off_t st_size; /* File size in bytes. */
#elif !defined(_FILE_OFFSET_BITS) || _FILE_OFFSET_BITS == 32
#if defined(__LITTLEENDIAN__)
 ino_t st_ino; /* File serial number. */
 ino_t st_ino_hi;
 off_t st_size;
 off_t st_size_hi;
#elif defined(__BIGENDIAN__)
 ino_t st_ino_hi;
 ino_t st_ino; /* File serial number. */
 off_t st_size_hi;
 off_t st_size;
#else
 #error endian not configured for system
#endif
#else
 #error _FILE_OFFSET_BITS value is unsupported
#endif
 dev_t st_dev; /* ID of the device containing the file. */
 dev_t st_rdev; /* Device ID. */
 uid_t st_uid; /* User ID of file. */
 gid_t st_gid; /* Group ID of file. */
 time_t st_mtime; /* Time of last data modification. */
 time_t st_atime; /* Time when file data was last accessed.*/
 time_t st_ctime; /* Time of last file status change. */
 mode_t st_mode; /* File types and permissions. */
 nlink_t st_nlink; /* Number of hard links to the file. */
 blksize_t st_blocksize; /* Size of a block used by st_nblocks. */
 _int32 st_nblocks; /* Number of blocks st_blocksize blocks. */
 blksize_t st_blksize; /* Preferred I/O block size for object. */
#if _FILE_OFFSET_BITS - 0 == 64
 blkcnt_t st_blocks; /* No. of 512-byte blocks allocated for a file. */
#elif !defined(_FILE_OFFSET_BITS) || _FILE_OFFSET_BITS == 32
#if defined(__LITTLEENDIAN__)
 blkcnt_t st_blocks; /* No. of 512-byte blocks allocated for a file. */
 blkcnt_t st_blocks_hi;
#elif defined(__BIGENDIAN__)
 blkcnt_t st_blocks_hi;
 blkcnt_t st_blocks;
#else
 #error endian not configured for system
#endif
#else
 #error _FILE_OFFSET_BITS value is unsupported
#endif
};

```



**Access permissions**

The access permissions for the file or directory are specified as a combination of bits in the *st\_mode* field of a **stat** structure. These bits are defined in `<sys/stat.h>`, and are described below:

| <b>Owner</b> | <b>Group</b> | <b>Others</b> | <b>Permission</b>                                                                                                                |
|--------------|--------------|---------------|----------------------------------------------------------------------------------------------------------------------------------|
| S_IRUSR      | S_IRGRP      | S_IROTH       | Read                                                                                                                             |
| S_IRWXU      | S_IRWXG      | S_IRWXO       | Read, write, execute/search. A bitwise inclusive OR of the other three constants. (S_IRWXU is OR of IRUSR, S_IWSUR and S_IXUSR.) |
| S_IWUSR      | S_IWGRP      | S_IWOTH       | Write                                                                                                                            |
| S_IXUSR      | S_IXGRP      | S_IXOTH       | Execute/search                                                                                                                   |

The following bits define miscellaneous permissions used by other implementations:

| <b>Bit</b> | <b>Equivalent</b> |
|------------|-------------------|
| S_IXEXEC   | S_IXUSR           |
| S_IREAD    | S_IRUSR           |
| S_IWRITE   | S_IWUSR           |

***st\_mode* bits**

The following bits are also encoded in the *st\_mode* field:

|         |                                                                                                                                                                                                  |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| S_ISUID | Set user ID on execution. The process's effective user ID is set to that of the owner of the file when the file is run as a program. On a regular file, this bit should be cleared on any write. |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

**S\_ISGID** Set group ID on execution. Set effective group ID on the process to the file's group when the file is run as a program. On a regular file, this bit should be cleared on any write.

**Macros**

The following symbolic names for the values of *st\_mode* are defined for these file types:

- S\_IFBLK** Block special.
- S\_IFCHR** Character special.
- S\_IFDIR** Directory.
- S\_IFIFO** FIFO special.
- S\_IFLNK** Symbolic link.
- S\_IFMT** Type of file.
- S\_IFNAM** Special named file.
- S\_IFREG** Regular.
- S\_IFSOCK** Socket.

The following macros test whether a file is of a specified type. The value *m* supplied to the macros is the value of the *st\_mode* field of a **stat** structure. The macros evaluate to a nonzero value if the test is true, and zero if the test is false.

- S\_ISBLK(*m*)** Test for block special file.
- S\_ISCHR(*m*)** Test for character special file.
- S\_ISDIR(*m*)** Test for directory file.
- S\_ISFIFO(*m*)** Test for FIFO.

|                    |                              |
|--------------------|------------------------------|
| <i>S_ISLNK(m)</i>  | Test for symbolic link.      |
| <i>S_ISNAM(m)</i>  | Test for special named file. |
| <i>S_ISREG(m)</i>  | Test for regular file.       |
| <i>S_ISSOCK(m)</i> | Test for socket.             |

These macros test whether a file is of the specified type. The value of the *buf* argument supplied to the macros is a pointer to a **stat** structure. The macro evaluates to a nonzero value if the specified object is implemented as a distinct file type and the specified file type is contained in the **stat** structure referenced by the pointer *buf*. Otherwise, the macro evaluates to zero.

*S\_TYPEISMQ(buf)*  
Test for message queue.

*S\_TYPEISSEM(buf)*  
Test for semaphore.

*S\_TYPEISSHM(buf)*  
Test for shared memory object.

These macros manipulate device IDs:

*major( device )*  
Extract the major number from a device ID.

*minor( device )*  
Extract the minor number from a device ID.

*makedev( node, major, minor)*  
Build a device ID from the given numbers. Currently, the *node* argument isn't used and must be zero.

The *st\_rdev* member of the **stat** structure is a device ID that consists of:

- a major number in the range 0 through 63
- a minor number in the range 0 through 1023.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

|              |                                                                                                                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of <i>path</i> .                                                                                                                        |
| EIO          | A physical error occurred on the block device.                                                                                                                                      |
| ELOOP        | Too many levels of symbolic links or prefixes.                                                                                                                                      |
| ENAMETOOLONG | The argument <i>path</i> exceeds PATH_MAX in length, or a pathname component is longer than NAME_MAX. These manifests are defined in the <code>&lt;limits.h&gt;</code> header file. |
| ENOENT       | The named file doesn't exist, or <i>path</i> is an empty string.                                                                                                                    |
| ENOSYS       | The <i>stat()</i> function isn't implemented for the filesystem specified in <i>path</i> .                                                                                          |
| ENOTDIR      | A component of <i>path</i> isn't a directory.                                                                                                                                       |
| EOVERFLOW    | The file size in bytes or the number of blocks allocated to the file or the file serial number can't be represented correctly in the structure pointed to by <i>buf</i> .           |

**Examples:**

Determine the size of a file:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>

int main(void)
{
 struct stat buf;

 if(stat("file", &buf) != -1) {
 printf("File size = %d\n", buf.st_size);
 }
 return EXIT_SUCCESS;
}
```

Determine the amount of free memory:

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>

int main () {
 struct stat buf;

 if (stat("/proc", &buf) == -1) {
 perror ("stat");
 return EXIT_FAILURE;
 } else {
 printf ("Free memory: %d bytes\n", buf.st_size);
 return EXIT_SUCCESS;
 }
}
```

**Classification:**

*stat()* is POSIX 1003.1; *stat64()* is for large-file support

**Safety**


---

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*errno, fstat(), fstat64(), lstat()*

## Synopsis:

```
#include <sys/statvfs.h>

int statvfs(const char *path,
 struct statvfs *buf);

int statvfs64(const char *path,
 struct statvfs64 *buf);
```

## Arguments:

*path* The name of a file that resides on the filesystem.

*buf* A pointer to a buffer where the function can store the information.

## Library:

libc

## Description:

The *statvfs()* function returns a “generic superblock” describing a filesystem; it can be used to acquire information about mounted filesystems. The *statvfs64()* function is a 64-bit version of *statvfs()*.

The filesystem type is known to the operating system. You don’t need to have read, write, or execute permission for the named file, but all directories listed in the path name leading to the file must be searchable.

The *buf* argument is a pointer to a **statvfs** or **statvfs64** structure that’s filled by the function. It contains at least:

**unsigned long** *f\_bsize*

The preferred filesystem blocksize.

**unsigned long** *f\_frsize*

The fundamental filesystem blocksize (if supported)

**fsblkcnt\_t** *f\_blocks*

The total number of blocks on the filesystem, in units of *f\_frsize*.

**fsblkcnt\_t** *f\_bfree*

The total number of free blocks.

**fsblkcnt\_t** *f\_bavail*

The number of free blocks available to a nonsuperuser.

**fsfilcnt\_t** *f\_files*

The total number of file nodes (inodes).

**fsfilcnt\_t** *f\_ffree*

The total number of free file nodes.

**fsfilcnt\_t** *f\_favail*

The number of inodes available to a nonsuperuser.

**unsigned long** *f\_fsid*

The filesystem ID (dev for now).

**char** *f\_basetype*[16]

The type of the target filesystem, as a null-terminated string.

**unsigned long** *f\_flag*

A bitmask of flags; the function can set these flags:

- **ST\_RDONLY** — read-only filesystem.
- **ST\_NOSUID** — the filesystem doesn't support **setuid/setgid** semantics.

**unsigned long** *f\_namemax*

The maximum filename length.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).



**Errors:**

|              |                                                                                                                                         |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied on a component of the path prefix.                                                                          |
| EFAULT       | The <i>path</i> or <i>buf</i> argument points to an illegal address.                                                                    |
| EINTR        | A signal was caught during execution.                                                                                                   |
| EIO          | An I/O error occurred while reading the filesystem.                                                                                     |
| ELOOP        | Too many symbolic links were encountered in translating <i>path</i> .                                                                   |
| EMULTIHOP    | Components of path require hopping to multiple remote machines and the filesystem type doesn't allow it.                                |
| ENAMETOOLONG | The length of a path component exceeds <b>{NAME_MAX}</b> characters, or the length of <i>path</i> exceeds <b>{PATH_MAX}</b> characters. |
| ENOENT       | Either a component of the path prefix or the file referred to by <i>path</i> doesn't exist.                                             |
| ENOLINK      | The <i>path</i> argument points to a remote machine and the link to that machine is no longer active.                                   |
| ENOTDIR      | A component of the path prefix of <i>path</i> isn't a directory.                                                                        |
| E_OVERFLOW   | One of the values to be returned can't be represented correctly in the structure pointed to by <i>buf</i> .                             |

**Classification:**

*statvfs()* is standard Unix; *statvfs64()* is for large-file support

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The values returned for *f\_files*, *f\_ffree*, and *f\_favail* might not be valid for NFS-mounted filesystems.

**See also:**

*chmod()*, *chown()*, *creat()*, *dup()*, *fcntl()*, *fstatvfs()*, *fstatvfs64()*, *link()*, *mknod()*, *open()*, *pipe()*, *read()*, *time()*, *unlink()*, *utime()*, *write()*

## Synopsis:

```
#include <stdio.h>

FILE * stderr;
```

## Description:

This global variable defines the standard error stream. It's set to the console by default, but you can redirect it by calling *freopen()*.

STDERR\_FILENO, which is defined in *<unistd.h>*, defines the file descriptor that corresponds to *stderr*.

## Classification:

ANSI

## See also:

*assert()*, *err()*, *errx()*, *getopt()*, *herror()*, *perror()*, *stdin*, *stdout*, *strerror()*, *verr()*, *verrx()*, *vwarn()*, *vwarnx()*, *warn()*, *warnx()*

## *stdin*

© 2004, QNX Software Systems Ltd.

*The standard input stream*

---

### **Synopsis:**

```
#include <stdio.h>

FILE * stdin;
```

### **Description:**

This global variable defines the standard input stream. It's set to the console by default, but you can redirect it by calling *freopen()*.

STDIN\_FILENO, which is defined in *<unistd.h>*, defines the file descriptor that corresponds to *stdin*.

### **Classification:**

ANSI

### **See also:**

*fgetchar()*, *getchar()*, *getchar\_unlocked()*, *gets()*, *getwchar()*, *scanf()*, *stderr*, *stdout*, *vscanf()*, *vwscanf()*

## Synopsis:

```
#include <stdio.h>

FILE * stdout;
```

## Description:

This global variable defines the standard output stream. It's set to the console by default, but you can redirect it by calling *freopen()*.

STDOUT\_FILENO, which is defined in *<unistd.h>*, defines the file descriptor that corresponds to *stdout*.

## Classification:

ANSI

## See also:

*fputc()*, *printf()*, *putchar()*, *putchar\_unlocked()*, *puts()*, *putwchar()*, *stderr*, *stdin*, *vprintf()*, *vwprintf()*, *wprintf()*

## ***straddstr()***

© 2004, QNX Software Systems Ltd.

*Concatenate one string on to the end of another*

### **Synopsis:**

```
#include <string.h>

int straddstr(const char * str,
 int len,
 char ** pbuf,
 size_t * pmaxbuf);
```

### **Arguments:**

|                |                                                                                                               |
|----------------|---------------------------------------------------------------------------------------------------------------|
| <i>str</i>     | The string that you want to add to the end of another.                                                        |
| <i>len</i>     | The number of characters from <i>str</i> that you want to add. If zero, the function adds all of <i>str</i> . |
| <i>pbuf</i>    | The address of a pointer to the destination buffer.                                                           |
| <i>pmaxbuf</i> | A pointer to the size of the destination buffer.                                                              |

### **Library:**

libc

### **Description:**

The *straddstr()* function adds *str* to the buffer pointed to by *pbuf*, respecting the maximum length indicated by *pmaxbuf*. The values of *pbuf* and *pmaxlen* are also updated.

### **Returns:**

The value of *len* if it's nonzero; otherwise, the length of *str* (i.e. *strlen(str)*).

### **Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strcat(), strcpy(), strncat(), strncpy()*

## ***strcasecmp()***

© 2004, QNX Software Systems Ltd.

*Compare two strings, ignoring case*

### **Synopsis:**

```
#include <strings.h>

int strcasecmp(const char* str1,
 const char* str2);
```

### **Arguments:**

*str1, str2*     The strings that you want to compare.

### **Library:**

libc

### **Description:**

The *strcasecmp()* function compares two strings, specified by *str1* and *str2*, ignoring the case of the characters.

### **Returns:**

< 0     *s1* is less than *s2*.  
0       *s1* is equal to *s2*.  
> 0     *s1* is greater than *s2*.

### **Examples:**

```
#include <stdio.h>
#include <strings.h>
#include <stdlib.h>

void compare(const char* s1, const char* s2)
{
 int retval;

 retval = strcasecmp(s1, s2);
 if(retval > 0) {
 printf("%s > %s\n", s1, s2);
 } else if(retval < 0) {
 printf("%s < %s\n", s1, s2);
 }
}
```



```
 } else {
 printf("%s == %s\n", s1, s2);
 }
}

int main(void)
{
 char* str1 = "abcdefg";
 char* str2 = "HIJ";
 char* str3 = "Abc";
 char* str4 = "aBCDEfg";

 compare(str1, str2);
 compare(str1, str3);
 compare(str1, str4);
 compare(str1, str1);

 compare(str2, str2);
 compare(str2, str3);
 compare(str2, str4);
 compare(str2, str1);

 return EXIT_SUCCESS;
}
```

This code produces output that looks like:

```
abcdefg < HIJ
abcdefg > Abc
abcdefg == aBCDEfg
abcdefg == abcdefg
HIJ == HIJ
HIJ > Abc
HIJ > aBCDEfg
HIJ > abcdefg
```

## Classification:

Standard Unix

### Safety

---

Cancellation point No

*continued...*

## **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

## **See also:**

*strcmp()*, *strcmpi()*, *strcoll()*, *stricmp()*, *strncasecmp()*, *strncmp()*,  
*strnicmp()*, *wscmp()*, *wscoll()*, *wcsncmp()*

## Synopsis:

```
#include <string.h>

char* strcat(char* dst,
 const char* src);
```

## Arguments:

*dst, src*     The strings that you want to concatenate.

## Library:

libc

## Description:

The *strcat()* function appends a copy of the string pointed to by *src* (including the terminating NUL character) to the end of the string pointed to by *dst*. The first character of *src* overwrites the NUL character at the end of *dst*.

## Returns:

The same pointer as *dst*.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];

 strcpy(buffer, "Hello ");
 strcat(buffer, "world");

 printf("%s\n", buffer);

 return EXIT_SUCCESS;
}
```

produces the output:

```
Hello world
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strncat()*, *strncpy()*, *strcpy()*

## Synopsis:

```
#include <string.h>

char* strchr(char* s,
 int c);
```

## Arguments:

- s*     The string that you want to search.
- c*     The character that you're looking for.

## Library:

libc

## Description:

The *strchr()* function finds the first occurrence of *c* (converted to a **char**) in the string pointed to by *s*. The terminating NUL character is considered to be part of the string.

## Returns:

A pointer to the located character, or NULL if *c* doesn't occur in the string.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];
 char* where;

 strcpy(buffer, "video x-rays");

 where = strchr(buffer, 'x');
```

```
 if(where == NULL) {
 printf("'x' not found\n");
 } else {
 printf("'x' found: %s\n", where);
 }

 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strcspn()*, *strpbrk()*, *strchr()*, *strspn()*, *strstr()*, *strtok()*,  
*strtok\_r()*, *wcchr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*

## Synopsis:

```
#include <string.h>

int strcmp(const char* s1,
 const char* s2);
```

## Arguments:

*s1*, *s2*     The strings that you want to compare.

## Library:

libc

## Description:

The *strcmp()* function compares the string pointed to by *s1* to the string pointed to by *s2*.

## Returns:

< 0     *s1* is less than *s2*.  
0       *s1* is equal to *s2*.  
> 0     *s1* is greater than *s2*.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 printf("%d\n", strcmp("abcdef", "abcdef"));
 printf("%d\n", strcmp("abcdef", "abc"));
 printf("%d\n", strcmp("abc", "abcdef"));
 printf("%d\n", strcmp("abcdef", "mnopqr"));
 printf("%d\n", strcmp("mnopqr", "abcdef"));

 return EXIT_SUCCESS;
}
```

produces the output:

```
0
1
-1
-1
1
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasecmp()*, *strcmpi()*, *strcoll()*, *stricmp()*, *strncasecmp()*, *strncmp()*, *strnicmp()*, *wscmp()*, *wscoll()*, *wcsncmp()*



**Synopsis:**

```
#include <string.h>

int strcmpi(const char* s1,
 const char* s2);
```

**Arguments:**

*s1*, *s2*     The strings that you want to compare.

**Library:**

libc

**Description:**

The *strcmpi()* function compares the string pointed to by *s1* to the string pointed to by *s2*, ignoring case.

All uppercase characters from *s1* and *s2* are mapped to lowercase for the purposes of doing the comparison. The *strcmpi()* function is identical to the *stricmp()* function.

**Returns:**

< 0     *s1* is less than *s2*.  
0       *s1* is equal to *s2*.  
> 0     *s1* is greater than *s2*.

**Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 printf("%d\n", strcmpi("AbCDEF", "abcdef"));
 printf("%d\n", strcmpi("abcdef", "ABC"));
 printf("%d\n", strcmpi("abc", "ABCdef"));
}
```

```
printf("%d\n", strcmpi("abcdef", "mnopqr"));
printf("%d\n", strcmpi("Mnopqr", "abcdef"));

return EXIT_SUCCESS;
}
```

produces the output:

```
0
100
-100
-12
12
```

## Classification:

QNX 4

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasecmp(), strcmp(), strcoll(), stricmp(), strncasecmp(), strncmp(), strnicmp(), wscmp(), wscoll(), wcsncmp()*

## Synopsis:

```
#include <string.h>

int strcoll(const char* s1,
 const char* s2);
```

## Arguments:

*s1*, *s2*     The strings that you want to compare.

## Library:

libc

## Description:

The *strcoll()* function compares the strings pointed to by *s1* and *s2*, using the collating sequence selected by the *setlocale()* function.

The *strcoll()* function is equivalent to *strcmp()* when the collating sequence is selected from the "C" locale.

## Returns:

< 0     *s1* is less than *s2*.  
0       *s1* is equal to *s2*.  
> 0     *s1* is greater than *s2*.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

char buffer[80] = "world";

int main(void)
{
 if(strcoll(buffer, "Hello") < 0) {
 printf("Less than\n");
 }
}
```

```
 }
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*setlocale(), strcasecmp(), strcmp(), strcmpi(), stricmp(), strncasecmp(), strncmp(), strnicmp(), wscmp(), wscoll(), wcsncmp()*

## Synopsis:

```
#include <string.h>

char* strcpy(char* dst,
 const char* src);
```

## Arguments:

*dst*     A pointer to where you want to copy the string.

*src*     The string that you want to copy.

## Library:

libc

## Description:

The *strcpy()* function copies the string pointed to by *src* (including the terminating NUL character) into the array pointed to by *dst*.



---

Copying of overlapping objects isn't guaranteed to work properly. See the *memmove()* function for information on copying objects that overlap.

---

## Returns:

The same pointer as *dst*.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char buffer[80];

 strcpy(buffer, "Hello ");
 strcat(buffer, "world");
```

```
 printf("%s\n", buffer);
 return EXIT_SUCCESS;
}
```

produces the output:

```
Hello world
```

## **Classification:**

ANSI

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## **See also:**

*memmove()*, *strdup()*, *strncpy()*, *wscpy()*, *wcsncpy()*, *wmemmove()*

*Count the characters at the beginning of a string that aren't in a given character set*

## Synopsis:

```
#include <string.h>

size_t strcspn(const char* str,
 const char* charset);
```

## Arguments:

*str*            The string that you want to search.

*charset*       The set of characters you want to look for.

## Library:

libc

## Description:

The *strcspn()* function finds the length of the initial segment of the string pointed to by *str* that consists entirely of characters *not* from the string pointed to by *charset*. The terminating NUL character isn't considered part of *str*.

## Returns:

The length of the initial segment.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 printf("%d\n", strcspn("abcbcadef", "cba"));
 printf("%d\n", strcspn("xxxbcadef", "cba"));
 printf("%d\n", strcspn("123456789", "cba"));

 return EXIT_SUCCESS;
}
```

produces the output:

0  
3  
9

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*, *strtok()*,  
*strtok\_r()*, *wchr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*



## Synopsis:

```
#include <string.h>

char* strdup(const char* src);
```

## Arguments:

*src*     The string that you want to copy.

## Library:

`libc`

## Description:

The *strdup()* function creates a duplicate of the string pointed to by *src*, and returns a pointer to the new copy.



---

The *strdup()* function allocates the memory for the new string by calling *malloc()*; it's up to you to release the memory by calling *free()*.

---

## Returns:

A pointer to a copy of the string, or NULL if an error occurred.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 char *dup;

 dup = strdup("Make a copy");
 printf("%s\n", dup);

 return EXIT_SUCCESS;
}
```

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*free()*, *malloc()*, *memmove()*, *strcpy()*, *strncpy()*, *wscpy()*, *wcsncpy()*,  
*wmemmove()*

## Synopsis:

```
#include <string.h>

char* strerror(int errnum);
```

## Arguments:

*errnum*     The error number that you want the message for. This function works for any valid *errno* value.

## Library:

libc

## Description:

The *strerror()* function maps the error number contained in *errnum* to an error message.

## Returns:

A pointer to the error message.



---

Don't modify the string that this function returns.

---

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <stdlib.h>

int main(void)
{
 FILE *fp;

 fp = fopen("file.name", "r");
 if(fp == NULL) {
 printf("Unable to open file: %s\n",
 strerror(errno));
 }
}
```

```
 return EXIT_SUCCESS;
 }
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno, perror(), stderr*

## Synopsis:

```
#include <time.h>

size_t strftime(char * s,
 size_t maxsize,
 const char * format,
 const struct tm * timeptr);
```

## Arguments:

|                |                                                                                    |
|----------------|------------------------------------------------------------------------------------|
| <i>s</i>       | A pointer to a buffer where the function can store the formatted time.             |
| <i>maxsize</i> | The maximum size of the buffer.                                                    |
| <i>format</i>  | The format that you want to use for the time; see “Formats,” below.                |
| <i>timeptr</i> | A pointer to a <b>tm</b> structure that contains the time that you want to format. |

## Library:

**libc**

## Description:

The *strftime()* function formats the time in the argument *timeptr* into the array pointed to by the argument *s*, according to the *format* argument.

## Formats

The *format* string consists of zero or more directives and ordinary characters. A directive consists of a % character followed by a character that determines the substitution that’s to take place. All ordinary characters are copied unchanged into the array. No more than *maxsize* characters are placed in the array.

Local timezone information is used as if from a call to *tzset()*.

|           |                                                                                              |
|-----------|----------------------------------------------------------------------------------------------|
| <b>%a</b> | Locale's abbreviated weekday name.                                                           |
| <b>%A</b> | Locale's full weekday name.                                                                  |
| <b>%b</b> | Locale's abbreviated month name.                                                             |
| <b>%B</b> | Locale's full month name.                                                                    |
| <b>%c</b> | Locale's appropriate date and time representation.                                           |
| <b>%d</b> | Day of the month as a decimal number (01-31).                                                |
| <b>%D</b> | Date in the format <i>mm/dd/yy</i> .                                                         |
| <b>%e</b> | Day of the month as a decimal number (1-31); single digits are preceded by a space.          |
| <b>%F</b> | The ISO standard date format; equivalent to <b>%Y-%m-%d</b> .                                |
| <b>%g</b> | The last two digits of the week-based year as a decimal number (00-99).                      |
| <b>%G</b> | The week-based year as a decimal number (e.g. 1998).                                         |
| <b>%h</b> | Locale's abbreviated month name.                                                             |
| <b>%H</b> | Hour (24-hour clock) as a decimal number (00-23).                                            |
| <b>%I</b> | Hour (12-hour clock) as a decimal number (01-12).                                            |
| <b>%j</b> | Day of the year as a decimal number (001-366).                                               |
| <b>%m</b> | Month as a decimal number (01-12).                                                           |
| <b>%M</b> | Minute as a decimal number (00-59).                                                          |
| <b>%n</b> | Newline character.                                                                           |
| <b>%p</b> | Locale's equivalent of either AM or PM.                                                      |
| <b>%r</b> | 12-hour clock time (01-12) using the AM/PM notation in the format <i>HH:MM:SS</i> (AM   PM). |
| <b>%R</b> | 24-hour notation; <b>%H:%M</b> .                                                             |

- %S** Second as a decimal number (00-59).
- %t** Tab character.
- %T** 24-hour clock time in the format *HH:MM:SS*.
- %u** Weekday as a decimal number (1-7), where Monday is 1.
- %U** Week number of the year as a decimal number (00-52), where Sunday is the first day of the week.
- %V** Week number of the year as a decimal number (01-53), where Monday is the first day of the week. The week containing January 1 is in the new year if four or more days are in the new year, otherwise it is the last week of the previous year.
- %w** Weekday as a decimal number (0-6), where 0 is Sunday.
- %W** Week number of the year as a decimal number (00-52), where Monday is the first day of the week.
- %x** Locale's appropriate date representation.
- %X** Locale's appropriate time representation.
- %Y** Year without century, as a decimal number (00-99).
- %Y** Year with century, as a decimal number.
- %z** Offset from UTC -0430 (4 hrs, 30 minutes behind UTC, west of Greenwich), or no characters if time zone isn't specified.
- %Z** Time zone name, or no characters if time zone isn't specified.
- %%** Character %.

Some of the above conversion specifiers can be modified with the prefix **E** or **O**. If alternative formats don't exist for the locale, they behave as if the unmodified conversion specifiers were called:

- %Ec** Alternative date and time representation.

|            |                                                                                                                                                             |
|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>%EC</b> | Alternative name of the the base year (period).                                                                                                             |
| <b>%Ex</b> | Alternative date representation.                                                                                                                            |
| <b>%EX</b> | Alternative time representation.                                                                                                                            |
| <b>%Ey</b> | Offset from <b>%EC</b> of the alternative year (only) representation.                                                                                       |
| <b>%EY</b> | Alternative year representation.                                                                                                                            |
| <b>%Od</b> | Day of the month using alternative numeric symbols.<br>Leading zeros are added if an alternative symbol for zero exists, otherwise leading spaces are used. |
| <b>%Oe</b> | Day of the month using alternative numeric symbols.<br>Leading spaces are used.                                                                             |
| <b>%OH</b> | 24-hour clock using alternative numeric symbols.                                                                                                            |
| <b>%OI</b> | 12-hour clock using alternative numeric symbols.                                                                                                            |
| <b>%Om</b> | Month using alternative numeric symbols.                                                                                                                    |
| <b>%OM</b> | Minutes using alternative numeric symbols.                                                                                                                  |
| <b>%OS</b> | Seconds using alternative numeric symbols.                                                                                                                  |
| <b>%Ou</b> | Alternative week day number representation (Monday=1).                                                                                                      |
| <b>%OU</b> | Alternative week day number representation (Rules correspond with <b>%U</b> ).                                                                              |
| <b>%OV</b> | Alternative week number representation. (Rules correspond with <b>%V</b> ).                                                                                 |
| <b>%Ow</b> | Weekday as a number using alternative numeric symbols (Sunday=0).                                                                                           |
| <b>%OW</b> | Week number of the year using alternative numeric symbols (Monday is the first day of the week).                                                            |
| <b>%Oy</b> | Year offset from <b>%C</b> using alternative numeric symbols.                                                                                               |



**Returns:**

The number of characters placed into the array, not including the terminating null character, or 0 if the number of characters exceeds *maxsize*; the string contents are indeterminate.

If an error occurs, *errno* is set.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

int main(void)
{
 time_t time_of_day;
 char buffer[80];

 time_of_day = time(NULL);
 strftime(buffer, 80, "Today is %A %B %d, %Y",
 localtime(&time_of_day));
 printf("%s\n", buffer);

 return EXIT_SUCCESS;
}
```

This produces the output:

```
Today is Thursday February 25, 1999
```

**Classification:**

ANSI

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*asctime()*, *asctime\_r()*, *ctime()*, *ctime\_r()*, *sprintf()*, **tm**, *tzset()*,  
*wcsftime()*

## Synopsis:

```
#include <string.h>

int stricmp(const char* s1,
 const char* s2);
```

## Arguments:

*s1*, *s2*     The strings that you want to compare.

## Library:

libc

## Description:

The *stricmp()* function compares, with case insensitivity, the string pointed to by *s1* to the string pointed to by *s2*. All uppercase characters from *s1* and *s2* are mapped to lowercase for the purposes of doing the comparison.

## Returns:

< 0     *s1* is less than *s2*.  
0       *s1* is equal to *s2*.  
> 0     *s1* is greater than *s2*.

## Examples:

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 printf("%d\n", stricmp("AbCDEF", "abcdef"));
 printf("%d\n", stricmp("abcdef", "ABC"));
 printf("%d\n", stricmp("abc", "ABCdef"));
 printf("%d\n", stricmp("Abcdef", "mnopqr"));
 printf("%d\n", stricmp("Mnopqr", "abcdef"));
}
```

```
 return EXIT_SUCCESS;
}
```

produces the output:

```
0
100
-100
-12
12
```

## Classification:

QNX 4

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasemp(), strcmp(), strcmpi(), strcoll(), strncasemp(), strncmp(), strnicmp(), wscmp(), wscoll(), wcsncmp()*

**Synopsis:**

```
#include <string.h>

size_t strlen(const char * s);
```

**Arguments:**

*s*     The string whose length you want to calculate.

**Library:**

libc

**Description:**

The *strlen()* function computes the length of the string pointed to by *s*.

**Returns:**

The number of characters that precede the terminating null character.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 printf("%d\n", strlen("Howdy"));
 printf("%d\n", strlen("Hello world\n"));
 printf("%d\n", strlen(""));

 return EXIT_SUCCESS;
}
```

produces the output:

```
5
12
0
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*wcslen()*

## Synopsis:

```
#include <string.h>

char* strlwr(char* s1);
```

## Arguments:

*s1*     The string that you want to convert to lowercase.

## Library:

libc

## Description:

The *strlwr()* function replaces the string *s1* with lowercase characters, by invoking *tolower()* for each character in the string.

## Returns:

The address of the string.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char source[] = { "A mixed-case STRING" };

int main(void)
{
 printf("%s\n", source);
 printf("%s\n", strlwr(source));
 printf("%s\n", source);
 return EXIT_SUCCESS;
}
```

produces the output:

```
A mixed-case STRING
a mixed-case string
A mixed-case string
```

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strupr(), tolower()*



### **Synopsis:**

```
#include <strings.h>

int strncasecmp(const char* str1,
 const char* str2,
 size_t n);
```

### **Arguments:**

*str1, str2*     The strings that you want to compare.

*n*             The maximum number of characters that you want to compare.

### **Library:**

libc

### **Description:**

The *strncasecmp()* function compares up to *n* characters in two strings, specified by *str1* and *str2*, ignoring the case of the characters.

### **Returns:**

< 0     *s1* is less than *s2*.

0       *s1* is equal to *s2*.

> 0     *s1* is greater than *s2*.

### **Examples:**

The following code:

```
#include <stdio.h>
#include <strings.h>
#include <stdlib.h>

void compare(const char *s1, const char *s2)
```

```
{
 int retval;

 retval = strncasecmp(s1, s2, 3);
 if(retval > 0) {
 printf("%s > %s\n", s1, s2);
 } else if(retval < 0) {
 printf("%s < %s\n", s1, s2);
 } else {
 printf("%s == %s\n", s1, s2);
 }
}

int main(void)
{
 char *str1 = "abcdefg";
 char *str2 = "HIJ";
 char *str3 = "Abc";
 char *str4 = "aBCDEfg";

 compare(str1, str2);
 compare(str1, str3);
 compare(str1, str4);
 compare(str1, str1);

 compare(str2, str2);
 compare(str2, str3);
 compare(str2, str4);
 compare(str2, str1);

 return EXIT_SUCCESS;
}
```

produces output that looks like:

```
abcdefg < HIJ
abcdefg == Abc
abcdefg == aBCDEfg
abcdefg == abcdefg
HIJ == HIJ
HIJ > Abc
HIJ > aBCDEfg
HIJ > abcdefg
```

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasecmp()*, *strcmp()*, *strcmpi()*, *strcoll()*, *stricmp()*, *strncmp()*,  
*strnicmp()*, *wscmp()*, *wscoll()*, *wcsncmp()*

## ***strncat()***

© 2004, QNX Software Systems Ltd.

*Concatenate two strings, up to a maximum length*

### **Synopsis:**

```
#include <string.h>

char* strncat(char* dst,
 const char* src,
 size_t n);
```

### **Arguments:**

*dst, src*     The strings that you want to concatenate.

*n*            The maximum number of characters that you want to add from the *src* string.

### **Library:**

`libc`

### **Description:**

The *strncat()* function appends no more than *n* characters of the string pointed to by *src* to the end of the string pointed to by *dst*. The first character of *src* overwrites the null character at the end of *dst*. This function always adds a terminating null character to the result.

### **Returns:**

The same pointer as *dst*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char buffer[80];

int main(void)
{
 strcpy(buffer, "Hello ");
 strncat(buffer, "world", 8);
 printf("%s\n", buffer);
}
```

```
 strncat(buffer, "*****", 4);
 printf("%s\n", buffer);
 return EXIT_SUCCESS;
}
```

produces the output:

```
Hello world
Hello world****
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcat()*

## ***strncmp()***

© 2004, QNX Software Systems Ltd.

*Compare two strings, up to a given length*

### **Synopsis:**

```
#include <string.h>

int strncmp(const char* s1,
 const char* s2,
 size_t n);
```

### **Arguments:**

*s1, s2*    The strings that you want to compare.

*n*        The maximum number of characters that you want to compare.

### **Library:**

libc

### **Description:**

The *strncmp()* function compares up to *n* characters from the strings pointed to by *s1* and *s2*.

### **Returns:**

< 0    *s1* is less than *s2*.

0      *s1* is equal to *s2*.

> 0    *s1* is greater than *s2*.

### **Examples:**

```
#include <stdio.h>
#include <string.h>
#include <stdlib.h>

int main(void)
{
 printf("%d\n", strncmp("abcdef", "abcDEF", 10));
 printf("%d\n", strncmp("abcdef", "abcDEF", 6));
}
```

```
printf("%d\n", strncmp("abcdef", "abcDEF", 3));
printf("%d\n", strncmp("abcdef", "abcDEF", 0));
return EXIT_SUCCESS;
}
```

produces the output:

```
1
1
0
0
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasemp(), strcmp(), strcmpi(), strcoll(), stricmp(), strncasemp(), strnicmp(), wscmp(), wcsoll(), wcsncmp()*

## ***strncpy()***

© 2004, QNX Software Systems Ltd.

*Copy a string, to a maximum length*

### **Synopsis:**

```
#include <string.h>

char* strncpy(char* dst,
 const char* src,
 size_t n);
```

### **Arguments:**

*dst*     A pointer to where you want to copy the string.

*src*     The string that you want to copy.

*n*       The maximum number of characters that you want to copy.

### **Library:**

`libc`

### **Description:**

The *strncpy()* function copies no more than *n* characters from the string pointed to by *src* into the array pointed to by *dst*.



---

Copying of overlapping objects isn't guaranteed to work properly. See the *memmove()* function if you wish to copy objects that overlap.

---

If the string pointed to by *src* is shorter than *n* characters, null characters are appended to the copy in the array pointed to by *dst*, until *n* characters in all have been written. If the string pointed to by *src* is longer than *n* characters, then the result isn't terminated by a null character.

### **Returns:**

The same pointer as *dst*.



**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 char buffer[15];

 printf("%s\n", strncpy(buffer, "abcdefg", 10));
 printf("%s\n", strncpy(buffer, "1234567", 6));
 printf("%s\n", strncpy(buffer, "abcdefg", 3));
 printf("%s\n", strncpy(buffer, "*****", 0));
 return EXIT_SUCCESS;
}
```

produces the output:

```
abcdefg
123456g
abc456g
abc456g
```

**Classification:**

ANSI

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memmove(), strcpy(), strdup(), wcsncpy(), wcsncpy(), wmemmove()*

## ***strnicmp()***

© 2004, QNX Software Systems Ltd.

*Compare two strings up to a given length, ignoring case*

### **Synopsis:**

```
#include <string.h>

int strnicmp(const char* s1,
 const char* s2,
 size_t len);
```

### **Arguments:**

*s1, s2*    The strings that you want to compare.

*len*        The maximum number of characters that you want to compare.

### **Library:**

libc

### **Description:**

The *strnicmp()* function compares up to *len* characters from the strings pointed to by *s1* and *s2*, ignoring case.

### **Returns:**

< 0    *s1* is less than *s2*.

0        *s1* is equal to *s2*.

> 0    *s1* is greater than *s2*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 printf("%d\n", strnicmp("abcdef", "ABCXXX", 10));
 printf("%d\n", strnicmp("abcdef", "ABCXXX", 6));
}
```

```
printf("%d\n", strnicmp("abcdef", "ABCXXX", 3));
printf("%d\n", strnicmp("abcdef", "ABCXXX", 0));
return EXIT_SUCCESS;
}
```

produces the output:

```
-20
-20
0
0
```

## Classification:

QNX 4

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strcasemp(), strcmp(), strcmpi(), strcoll(), stricmp(), strncasemp(), strncmp(), wcsmp(), wcsoll(), wcsncmp()*

## ***strnset()***

© 2004, QNX Software Systems Ltd.

*Fill a string with a given character, to a given length*

### **Synopsis:**

```
#include <string.h>

char * strnset(char * s1,
 int fill,
 size_t len);
```

### **Arguments:**

*s1*     The string that you want to fill.

*fill*    The value that you want to fill the string with.

*len*     The number of characters to fill.

### **Library:**

`libc`

### **Description:**

The *strnset()* function fills the string *s1* with the value of the argument *fill*, converted to be a character value. When the value of *len* is greater than the length of the string, the entire string is filled. Otherwise, that number of characters at the start of the string are set to the fill character.

### **Returns:**

The address of the string, *s1*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char source[] = { "A sample STRING" };

int main(void)
{
 printf("%s\n", source);
}
```

```
 printf("%s\n", strnset(source, '=', 100));
 printf("%s\n", strnset(source, '*', 7));
 return EXIT_SUCCESS;
}
```

produces the output:

```
A sample STRING
=====
*****=====
```

## Classification:

QNX 4

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strset()*

## ***strpbrk()***

© 2004, QNX Software Systems Ltd.

*Find the first character in a string that's in a given character set*

### **Synopsis:**

```
#include <string.h>

char* strpbrk(char* str,
 char* charset);
```

### **Arguments:**

*str*            The string that you want to search.

*charset*       The set of characters you want to look for.

### **Library:**

libc

### **Description:**

The *strpbrk()* function locates the first occurrence in the string pointed to by *str* of any character from the string pointed to by *charset*.

### **Returns:**

A pointer to the located character, or NULL if no character from *charset* occurs in *str*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 char* p = "Find all vowels";

 while(p != NULL) {
 printf("%s\n", p);
 p = strpbrk(p+1, "aeiouAEIOU");
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
Find all vowels
ind all vowels
all vowels
owels
els
```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memchr()*, *strchr()*, *strcspn()*, *strchr()*, *strspn()*, *strstr()*, *strtok()*,  
*strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*

## ***strrchr()***

© 2004, QNX Software Systems Ltd.

*Find the last occurrence of a character in a string*

### **Synopsis:**

```
#include <string.h>

const char* strrchr(const char* s,
 int c);
```

### **Arguments:**

- s*     The string that you want to search.
- c*     The character that you're looking for.

### **Library:**

`libc`

### **Description:**

The *strrchr()* function locates the last occurrence of *c* (converted to a **char**) in the string pointed to by *s*. The terminating null character is considered to be part of the string.

### **Returns:**

A pointer to the located character, or a NULL pointer if the character doesn't occur in the string.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 printf("%s\n", strrchr("abcdeabcde", 'a'));
 if(strrchr("abcdeabcde", 'x') == NULL)
 printf("NULL\n");
 return EXIT_SUCCESS;
}
```

produces the output:



abcde  
NULL

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strspn()*, *strstr()*, *strtok()*,  
*strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*

## ***strrev()***

© 2004, QNX Software Systems Ltd.

*Reverse a string*

---

### **Synopsis:**

```
#include <string.h>

char* strrev(char* s1);
```

### **Arguments:**

*s1*     The string that you want to reverse.

### **Library:**

libc

### **Description:**

The *strrev()* function replaces the string *s1* with a string whose characters are in the reverse order.

### **Returns:**

The address of the string, *s1*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char source[] = { "A sample STRING" };

int main(void)
{
 printf("%s\n", source);
 printf("%s\n", strrev(source));
 printf("%s\n", strrev(source));
 return EXIT_SUCCESS;
}
```

produces the output:

```
A sample STRING
GNIRTS elpmas A
A sample STRING
```

## Classification:

QNX 4

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## ***strsep()***

© 2004, QNX Software Systems Ltd.

*Separate a string into pieces marked by given delimiters*

### **Synopsis:**

```
#include <string.h>

char *strsep(char **stringp,
 char *delim);
```

### **Arguments:**

*stringp*     The address of a pointer to the string that you want to break into pieces; see below.

*delim*       A set of characters that delimit the pieces in the string.

### **Library:**

libc

### **Description:**

The *strsep()* function looks in the null-terminated string pointed to by *stringp* for the first occurrence of any character in *delim* and replaces this with a `\0`, records the location of the next character in *stringp*, then returns the original value of *stringp*. If no delimiter characters are found, *strsep()* sets *stringp* to NULL; if *stringp* is initially NULL, *strsep()* returns NULL.

### **Returns:**

A pointer to the original value of *stringp*.

### **Examples:**

Parse strings containing runs of whitespace, making up an argument vector:

```
char inputstring[100];
char **argv[51], **ap = argv, *p, *val;

/* set up inputstring */
for (p = inputstring; p != NULL;) {
```

```
 while ((val = strsep(&p, " \t")) != NULL && *val == '\0');
 *ap++ = val;
 }
 *ap = 0;
```

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strtok()*, *strtok\_r()*, *wcstok()*

## ***strset()***

© 2004, QNX Software Systems Ltd.

*Fill a string with a given character*

### **Synopsis:**

```
#include <string.h>

char* strset(char* s1,
 int fill);
```

### **Arguments:**

*s1*     The string that you want to fill.  
*fill*    The value that you want to fill the string with.

### **Library:**

libc

### **Description:**

The *strset()* function fills the string pointed to by *s1* with the character *fill*. The terminating null character in the original string remains unchanged.

### **Returns:**

The address of the string, *s1*.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char source[] = { "A sample STRING" };

int main(void)
{
 printf("%s\n", source);
 printf("%s\n", strset(source, '='));
 printf("%s\n", strset(source, '*'));
 return EXIT_SUCCESS;
}
```

produces the output:

A sample STRING  
=====  
\*\*\*\*\*

## Classification:

QNX 4

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strnset()*

## ***strsignal()***

© 2004, QNX Software Systems Ltd.

*Return the description of a signal*

---

### **Synopsis:**

```
#include <string.h>

char *strsignal(int signo);
```

### **Arguments:**

*signo*     The signal number that you want the description of.

### **Library:**

`libc`

### **Description:**

The *strsignal()* function returns a pointer to the language-dependent string describing a signal.

### **Returns:**

A pointer to the description of the signal, or NULL if *signo* isn't a valid signal number. This array will be overwritten by subsequent calls to *strsignal()*.



---

Don't modify the array returned by this function.

---

### **Classification:**

Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*setlocale()*

## ***strspn()***

© 2004, QNX Software Systems Ltd.

*Count the characters at the beginning of a string that are in a given character set*

### **Synopsis:**

```
#include <string.h>

size_t strspn(const char* str,
 const char* charset);
```

### **Arguments:**

*str*            The string that you want to search.

*charset*       The set of characters you want to look for.

### **Library:**

libc

### **Description:**

The *strspn()* function computes the length of the initial segment of the string pointed to by *str* that consists of characters from the string pointed to by *charset*. The terminating null character isn't considered to be part of *charset*.

### **Returns:**

The length of the segment.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 printf("%d\n", strspn("out to lunch", "aeiou"));
 printf("%d\n", strspn("out to lunch", "xyz"));
 return EXIT_SUCCESS;
}
```

produces the output:

2  
0

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strstr()*, *strtok()*,  
*strtok\_r()*, *wcchr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*

# ***strstr()***

© 2004, QNX Software Systems Ltd.

*Find one string inside another*

## **Synopsis:**

```
#include <string.h>

char* strstr(char* str,
 char* substr);
```

## **Arguments:**

*str*            The string that you want to search.  
*substr*        The string that you're looking for.

## **Library:**

libc

## **Description:**

The *strstr()* function locates the first occurrence in the string pointed to by *str* of the sequence of characters (excluding the terminating null character) in the string pointed to by *substr*.

## **Returns:**

A pointer to the located string, or NULL if the string isn't found.

## **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 printf("%s\n", strstr("This is an example", "is"));
 return EXIT_SUCCESS;
}
```

produces the output:

```
is is an example
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strtok()*,  
*strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcsstr()*, *wcstok()*

## ***strtod()***

© 2004, QNX Software Systems Ltd.

*Convert a string into a double*

### **Synopsis:**

```
#include <stdlib.h>

double strtod(const char *ptr,
 char **endptr);
```

### **Arguments:**

*ptr*            A pointer to the string to parse.

*endptr*        If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.

### **Library:**

`libc`

### **Description:**

The *strtod()* function converts the string pointed to by *ptr* to *double* representation. The function recognizes a string containing the following:

- optional white space
- an optional plus or minus sign
- a sequence of digits containing an optional decimal point
- an optional **e** or **E**, followed by an optionally signed sequence of digits.

The conversion ends at the first unrecognized character. If *endptr* isn't NULL, a pointer to the unrecognized character is stored in the object *endptr* points to.

## Returns:

The converted value. If the correct value would cause overflow, plus or minus `HUGE_VAL` is returned according to the sign, and *errno* is set to `ERANGE`. If the correct value would cause underflow, then zero is returned, and *errno* is set to `ERANGE`.

This function returns zero when the input string can't be converted. If an error occurs, *errno* indicates the error detected.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 double pi;

 pi = strtod("3.141592653589793", NULL);
 printf("pi=%17.15f\n", pi);
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atof(), errno*



### **Synopsis:**

```
#include <inttypes.h>

intmax_t strtoimax (const char * nptr,
 char ** endptr,
 int base);

uintmax_t strtoumax (const char * nptr,
 char ** endptr,
 int base);
```

### **Arguments:**

- nptr*        A pointer to the string to parse.
- endptr*     If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*        The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.

**Library:**

`libc`

**Description:**

The *strtoimax()* and *strtoumax()* functions are the same as the *strtol()*, *strtoll()*, *strtoul()*, and *strtoull()* functions except that they return objects of type `intmax_t` and `uintmax_t`.

**Returns:**

The converted value.

If the correct value causes an overflow, `INTMAX_MAX` | `UINTMAX_MAX` or `INTMAX_MIN` is returned according to the sign and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EINVAL`.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strtol()*, *strtoul()*

**Synopsis:**

```
#include <string.h>

char* strtok(char* s1,
 const char* s2);
```

**Arguments:**

- s1*     NULL, or the string that you want to break into tokens; see below.
- s2*     A set of the characters that separate the tokens.

**Library:**

`libc`

**Description:**

The *strtok()* function breaks the string pointed to by *s1* into a sequence of tokens, each of which is delimited by a character from the string pointed to by *s2*.

The first call to *strtok()* returns a pointer to the first token in the string pointed to by *s1*. Subsequent calls to *strtok()* must pass a NULL pointer as the first argument, in order to get the next token in the string. The set of delimiters used in each of these calls to *strtok()* can be different from one call to the next.

The first call in the sequence searches *s1* for the first character that isn't contained in the current delimiter string *s2*. If no such character is found, then there are no tokens in *s1*, and *strtok()* returns a NULL pointer. If such a character is found, it's the start of the first token.

The *strtok()* function then searches from there for a character that's contained in the current delimiter string. If no such character is found, the current token extends to the end of the string pointed to by *s1*. If such a character is found, it's overwritten by a null character, which terminates the current token. The *strtok()* function saves a pointer to

the following character, from which the next search for a token will start when the first argument is a NULL pointer.



---

You might want to keep a copy of the original string because *strtok()* is likely to modify it.

---

## Returns:

A pointer to the token found, or NULL if no token was found.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

int main(void)
{
 char* p;
 char* buffer;
 char* delims = { " .," };

 buffer = strdup("Find words, all of them.");
 printf("%s\n", buffer);
 p = strtok(buffer, delims);
 while(p != NULL) {
 printf("word: %s\n", p);
 p = strtok(NULL, delims);
 }
 printf("%s\n", buffer);
 return EXIT_SUCCESS;
}
```

produces the output:

```
Find words, all of them.
word: Find
word: words
word: all
word: of
word: them
Find
```

## Classification:

ANSI

### Safety

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## See also:

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strset()*, *strspn()*,  
*strstr()*, *strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*,  
*wcsspn()*, *wcsstr()*, *wcstok()*

## ***strtok\_r()***

© 2004, QNX Software Systems Ltd.

*Break a string into tokens (reentrant)*

### **Synopsis:**

```
#include <string.h>

char* strtok_r(char* s,
 const char* sep,
 char** lasts);
```

### **Arguments:**

- s1*      NULL, or the string that you want to break into tokens; see below.
- s2*      A set of the characters that separate the tokens.
- lasts*    The address of a pointer to a character, which the function can use to store information necessary for it to continue scanning the same string.

### **Library:**

`libc`

### **Description:**

The function *strtok\_r()* breaks the string *s* into a sequence of tokens, each of which is delimited by a character from the string pointed to by *sep*.

In the first call to *strtok\_r()*, *s* must point to a null-terminated string, *sep* points to a null-terminated string of separator characters, and *lasts* is ignored. The *strtok\_r()* function returns a pointer to the first character of the first token, writes a NULL character into *s* immediately following the returned token, and updates *lasts*.

In subsequent calls, *s* must be a NULL pointer and *lasts* must be unchanged from the previous call so that subsequent calls will move through the string *s*, returning successive tokens until no tokens remain. The separator string *sep* may be different from call to call. When no tokens remain in *s*, a NULL pointer is returned.

**Returns:**

A pointer to the token found, or NULL if no token was found.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strset()*, *strspn()*,  
*strstr()*, *strtok()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcssstr()*, *wcstok()*

## ***strtol()*, *strtoll()***

© 2004, QNX Software Systems Ltd.

*Convert a string into a long integer*

### **Synopsis:**

```
#include <stdlib.h>

long int strtol(const char * ptr,
 char ** endptr,
 int base);

int64_t strtoll(const char * ptr,
 char ** endptr,
 int base);
```

### **Arguments:**

- ptr*            A pointer to the string to parse.
- endptr*        If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*           The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.



**Library:**`libc`**Description:**

The *strtol()* function converts the string pointed to by *ptr* to an object of type `long int`; *strtoll()* converts the string pointed to by *ptr* to an object of type `int64_t` (`long long`).

These functions recognize strings that contain the following:

- optional white space
- an optional plus or minus sign
- a sequence of digits and letters.

The conversion ends at the first unrecognized character. If *endptr* isn't NULL, a pointer to the unrecognized character is stored in the object *endptr* points to.

**Returns:**

The converted value.

If the correct value causes an overflow, `LONG_MAX` | `LONGLONG_MAX` or `LONG_MIN` | `LONGLONG_MIN` is returned according to the sign, and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EDOM`.

**Examples:**

```
#include <stdlib.h>

int main(void)
{
 long int v;

 v = strtol("12345678", NULL, 10);
 return EXIT_SUCCESS;
}
```

**Classification:**

*strtol()* is ANSI; *strtoll()* is Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atoi()*, *atol()*, *errno*, *itoa()*, *ltoa()*, *sscanf()*, *strtoul()*, *ultoa()*, *utoa()*

### **Synopsis:**

```
#include <stdlib.h>

unsigned long int strtoul(const char * ptr,
 char ** endptr,
 int base);

uint64_t strtoull(const char * ptr,
 char ** endptr,
 int base);
```

### **Arguments:**

- ptr*            A pointer to the string to parse.
- endptr*        If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*           The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.

## Library:

`libc`

## Description:

The *strtoul()* function converts the string pointed to by *ptr* to an **unsigned long**; *strtoull()* converts the string pointed to by *ptr* to a **uint64\_t (unsigned long long)**.

These functions recognize strings that contain the following:

- optional white space
- a sequence of digits and letters.

The conversion ends at the first unrecognized character. A pointer to that character is stored in the object *endptr* points to, if *endptr* isn't NULL.

## Returns:

The converted value.

If the correct value causes an overflow, `ULONG_MAX | ULLONG_MAX` is returned and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EDOM`.

## Examples:

```
#include <stdlib.h>

int main(void)
{
 unsigned long int v;

 v = strtoul("12345678", NULL, 10);
 return EXIT_SUCCESS;
}
```

**Classification:**

*strtoul()* is ANSI; *strtoull()* is Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atoi()*, *atol()*, *errno*, *itoa()*, *ltoa()*, *sscanf()*, *strtol()*, *ultoa()*, *utoa()*

## ***strupr()***

© 2004, QNX Software Systems Ltd.

*Convert a string to uppercase*

### **Synopsis:**

```
#include <string.h>

char* strupr(char* s1);
```

### **Arguments:**

*s1*     The string that you want to convert to uppercase.

### **Library:**

libc

### **Description:**

The *strupr()* function replaces the string *s1* with uppercase characters, by invoking *toupper()* for each character in the string.

### **Returns:**

The address of the string.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>

char source[] = { "A mixed-case STRING" };

int main(void)
{
 printf("%s\n", source);
 printf("%s\n", strupr(source));
 printf("%s\n", source);
 return EXIT_SUCCESS;
}
```

produces the output:

```
A mixed-case STRING
A MIXED-CASE STRING
A MIXED-CASE STRING
```

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*strlwr()*, *toupper()*

## ***strxfrm()***

© 2004, QNX Software Systems Ltd.

*Transform one string into another, to a given length*

### **Synopsis:**

```
#include <string.h>

size_t strxfrm(char* dst,
 const char* src,
 size_t n);
```

### **Arguments:**

*dst*    The string that you want to transform.

*src*    The string that you want to place in *dst*.

*n*      The maximum number of characters to transform.

### **Library:**

libc

### **Description:**

The *strxfrm()* function transforms, for no more than *n* characters, the string pointed to by *src* to the buffer pointed to by *dst*. The transformation uses the collating sequence selected by *setlocale()* so that two transformed strings compare identically (using the *strncmp()* function) to a comparison of the original two strings using *strcoll()*.

If the collating sequence is selected from the "C" locale, *strxfrm()* is equivalent to *strncpy()*, except that *strxfrm()* doesn't pad the *dst* argument with null characters when the argument *src* is shorter than *n* characters.

### **Returns:**

The length of the transformed string. If this length is more than *n*, the contents of the array pointed to by *dst* are indeterminate.






---

If an error occurs, *strxfrm()* sets *errno* and returns 0. Since the function could also return zero on success, the only way to tell that an error has occurred is to set *errno* to 0 before calling *strxfrm()* and check *errno* afterward.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <locale.h>

char src[] = { "A sample STRING" };
char dst[20];

int main(void)
{
 size_t len;

 setlocale(LC_ALL, "C");
 printf("%s\n", src);
 len = strxfrm(dst, src, 20);
 printf("%s (%u)\n", dst, len);
 return EXIT_SUCCESS;
}
```

produces the output:

```
A sample STRING
A sample STRING (15)
```

## Classification:

ANSI

### Safety

---

Cancellation point    No

Interrupt handler      Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*setlocale(), strcoll(), wcsxfrm()*

## Synopsis:

```
#include <unistd.h>

void swab(const void * src,
 void * dest,
 ssize_t nbytes);
```

## Arguments:

*src*            A pointer to the buffer that you want to copy the bytes from.

*dest*           A pointer to the buffer where you want the function to copy the bytes.

*nbytes*         The number of bytes that you want to copy and swap.

## Library:

`libc`

## Description:

The *swab()* function copies *nbytes* bytes, pointed to by *src*, to the object pointed to by *dest*, exchanging adjacent bytes. The *nbytes* argument should be even.

| <b>If <i>nbytes</i> is:</b> | <b>Then:</b>                                                                                      |
|-----------------------------|---------------------------------------------------------------------------------------------------|
| Odd                         | <i>nbytes</i> -1 bytes are copied and exchanged. The disposition of the last byte is unspecified. |
| Negative                    | <i>swab()</i> does nothing.                                                                       |

If copying takes place between objects that overlap, the behavior is undefined.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ENDIAN\_SWAP32()*, *ENDIAN\_SWAP64()*

## Synopsis:

```
#include <wchar.h>

int swprintf(wchar_t * ws,
 size_t n,
 const wchar_t * format,
 ...);
```

## Arguments:

*ws*            A pointer to the buffer where you want to function to store the formatted string.

*n*             The maximum number of wide characters to store in the buffer, including a terminating null character.

*format*        A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

## Library:

`libc`

## Description:

The *swprintf()* function is similar to *fwprintf()* except that *swprintf()* places the generated output into the wide-character array pointed to by *buf*, instead of writing it to a file. A null character is placed at the end of the generated character string.

The *swprintf()* function is the wide-character version of *sprintf()*.

## Returns:

The number of wide characters written, excluding the terminating **NUL**, or a negative number if an error occurred (*errno* is set).

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno, fprintf(), fwprintf(), printf(), snprintf(), sprintf(), vfprintf(),  
vfwprintf(), vprintf(), vsnprintf(), vsprintf(), vswprintf(), vwprintf(),  
wprintf()*

## Synopsis:

```
#include <wchar.h>

int swscanf(const wchar_t * ws,
 const wchar_t * format,
 ...);
```

## Arguments:

*ws*            The wide-character string that you want to read from.

*format*        A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

## Library:

`libc`

## Description:

The *swscanf()* function scans input from the wide-character string *ws*, under control of the argument *format*. Following the format string is the list of addresses of items to receive values.

The *swscanf()* function is the wide-character version of *sscanf()*.

## Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF when the scanning is terminated by reaching the end of the input string.

## Classification:

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*, *fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *vfscanf()*, *vwscanf()*,  
*vscanf()*, *vsscanf()*, *vswscanf()*, *vscanf()*, *wscanf()*



## Synopsis:

```
#include <unistd.h>

int symlink(const char* pname,
 const char* slink);
```

## Arguments:

*pname*     The path that you want to link to.

*slink*     The symbolic link that you want to create.

## Library:

`libc`

## Description:

The *symlink()* function creates a symbolic link named *slink* that contains the pathname specified by *pname* (*slink* is the name of the symbolic link created, *pname* is the pathname contained in the symbolic link).

File access checking isn't performed on the file named by *pname*, and the file need not exist.

If the *symlink()* function is unsuccessful, any file named by *slink* is unaffected.

## Returns:

0     Success.

-1    An error occurred (*errno* is set).

## Errors:

EACCES     A component of the *slink* path prefix denies search permission, or write permission is denied in the parent directory of the symbolic link to be created.

|              |                                                                                                                                              |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| EEXIST       | A file named by <i>slink</i> already exists.                                                                                                 |
| ELOOP        | A loop exists in symbolic links encountered during resolution of the <i>slink</i> argument, and it resolves to more than SYMLOOP_MAX levels. |
| ENAMETOOLONG | A component of the path specified by <i>slink</i> exceeds NAME_MAX bytes, or the length of the entire pathname exceeded PATH_MAX characters. |
| ENOSPC       | The new symbolic link can't be created because there's no space left on the filesystem that will contain the symbolic link.                  |
| ENOSYS       | The <i>symlink()</i> function isn't implemented for the filesystem specified in <i>slink</i> .                                               |
| ENOTDIR      | A component of the path prefix of <i>slink</i> isn't a directory.                                                                            |
| EROFS        | The file <i>slink</i> would reside on a read-only filesystem.                                                                                |

**Examples:**

```
/*
 * create a symbolic link to "/usr/nto/include"
 */
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(void)
{
 if(symlink("/usr/nto/include", "slink") == -1) {
 perror("slink -> /usr/nto/include");
 exit(EXIT_FAILURE);
 }
 exit(EXIT_SUCCESS);
}
```

## Classification:

POSIX 1003.1a

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *link()*, *lstat()*, *pathmgr\_symlink()*, *pathmgr\_unlink()*, *readlink()*, *unlink()*

# ***sync()***

© 2004, QNX Software Systems Ltd.

*Synchronize filesystem updates*

## **Synopsis:**

```
#include <unistd.h>

void sync(void);
```

## **Library:**

libc

## **Description:**

The *sync()* function queues all the modified block buffers for writing, and returns; it doesn't wait for the actual I/O to take place. Use this function — or *fsync()* for a single file — to ensure consistency of the entire on-disk filesystem with the contents of the in-memory buffer cache.

## **Returns:**

-1 on error; any other value on success.

## **Classification:**

Standard Unix

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*fdatsync()*, *fsync()*

## **SyncCondvarSignal(), SyncCondvarSignal\_r()** © 2004,

QNX Software Systems Ltd.

*Wake up any threads that are blocked on a synchronization object*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SyncCondvarSignal(sync_t* sync,
 int broadcast);

int SyncCondvarSignal_r(sync_t* sync,
 int broadcast);
```

### **Arguments:**

|                  |                                                                                                                                                                                                                          |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>sync</i>      | A pointer to a <b>sync_t</b> for the synchronization object. You must have initialized this argument by calling <i>SyncTypeCreate()</i> or statically initialized it with the manifest <b>PTHREAD_COND_INITIALIZER</b> . |
| <i>broadcast</i> | Zero if you want to make ready to run the thread with the highest priority that's been waiting the longest, or nonzero if you want to make all waiting threads ready.                                                    |

### **Library:**

**libc**

### **Description:**

The *SyncCondvarSignal()* and *SyncCondvarSignal\_r()* kernel calls wake up one or all threads that are blocked on the synchronization object *sync*.

These functions are similar, except in the way they indicate errors. See the Returns section for details.

In all cases, each awakened thread attempts to reacquire the controlling mutex passed in *SyncCondvarWait()* before control is returned to the thread. If the mutex is already locked when the kernel attempts to lock it, the thread becomes blocked on the mutex until it's unlocked.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *SyncCondvarSignal()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

### *SyncCondvarSignal\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

- EFAULT     A fault occurred when the kernel tried to access *sync*.
- EINVAL     The synchronization ID specified in *sync* doesn't exist.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

# **SyncCondvarSignal(), SyncCondvarSignal\_r()** © 2004,

QNX Software Systems Ltd.

---

## **See also:**

*pthread\_cond\_broadcast(), pthread\_cond\_signal(),  
pthread\_cond\_wait(), SyncCondvarWait()*



## ***SyncCondvarWait(), SyncCondvarWait\_r()***

*Block a thread on a synchronization object*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SyncCondvarWait(sync_t * sync,
 sync_t * mutex);

int SyncCondvarWait_r(sync_t * sync,
 sync_t * mutex);
```

### **Arguments:**

*sync*      A pointer to a **sync\_t** for the synchronization object. You must have initialized this argument by calling *SyncTypeCreate()* or statically initialized it with the manifest `PTHREAD_COND_INITIALIZER`.

*mutex*     The mutex that's associated with the condition variable. You must lock this mutex by calling *SyncMutexLock()* (or the POSIX *pthread\_mutex\_lock()* cover routine). The kernel releases the mutex lock in the kernel when it blocks the thread on *sync*.

### **Library:**

`libc`

### **Description:**

The *SyncCondvarWait()* and *SyncCondvarWait\_r()* kernel calls block the calling thread on the synchronization object, *sync*. If more than one thread is blocked on the object, they're queued in priority order.

These functions are similar, except in the way they indicate errors. See the Returns section for details.

The blocked thread can be unblocked by any one of the following conditions:

## Condition variable signalled

The condition variable was signaled by a call to *SyncCondvarSignal()*, that determined that this thread should be awakened.

Before returning from *SyncCondvarWait()*, *mutex* is reacquired. If *mutex* is locked, the thread enters into the STATE\_MUTEX state waiting for *mutex* to be unlocked. At this point it's as though you had called *SyncMutexLock(mutex)*.

## Timeout

The wait was terminated by a timeout initiated by a previous call to *TimerTimeout()*.

Before returning from *SyncCondvarWait()*, *mutex* is reacquired. If *mutex* is locked, the thread enters into the STATE\_MUTEX state waiting for *mutex* to be unlocked. At this point it's as though you had called *SyncMutexLock(mutex)*.

## POSIX signal

The wait was terminated by an unmasked signal initiated by a call to *SignalKill()*. If a signal handler has been set up, the signal handler runs with *mutex* unlocked. On return from the signal handler, *mutex* is reacquired. If *mutex* is locked, the thread enters into the STATE\_MUTEX state waiting for *mutex* to be unlocked. At this point, it's as though you had called *SyncMutexLock(mutex)*.

## Thread cancellation

The wait was terminated by a thread cancellation initiated by a call to *ThreadCancel()*. Before calling the cancellation handler, *mutex* is reacquired. If *mutex* is locked, the thread enters into the STATE\_MUTEX state waiting for *mutex* to be unlocked. At this point, it's as though you had called *SyncMutexLock(mutex)*.

In all cases, *mutex* is reacquired before the call returns. If the thread enters the STATE\_MUTEX state, the rules governing *SyncMutexLock()* are in effect.

Condition variables are used to block a thread until a certain condition is satisfied. Spurious wakeups may occur due to timeouts, signals, and broadcast condition variable signals. Therefore, you should always reevaluate the condition, even on a successful return. The easiest way to do this is with a while loop. For example:

```
SyncMutexLock (&mutex);

while (some_condition) {
 SyncCondvarWait (&condvar, &mutex);
}

SyncMutexUnlock (&mutex);
```

### Blocking states

#### STATE\_CONDVAR

The calling thread blocks waiting for the condition variable to be signaled.

#### STATE\_MUTEX

The thread was unblocked from the STATE\_CONDVAR state and while trying to reacquire the controlling mutex, found the mutex was locked by another thread.

### Returns:

The only difference between these functions is the way they indicate errors:

#### *SyncCondvarWait()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

# SyncCondvarWait(), SyncCondvarWait\_r()

© 2004, QNX

Software Systems Ltd.

---

## *SyncCondvarWait\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function returns any value in the Errors section.

### Errors:

|           |                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------|
| EAGAIN    | On the first use of a statically initialized <i>sync</i> , all kernel synchronization objects were in use. |
| EFAULT    | A fault occurred when the kernel tried to access <i>sync</i> or <i>mutex</i> .                             |
| EINVAL    | The synchronization ID specified in <i>sync</i> doesn't exist.                                             |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                           |

### Classification:

QNX Neutrino

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_cond\_broadcast()*, *pthread\_cond\_signal()*,  
*pthread\_cond\_wait()*, *pthread\_mutex\_lock()*, *SignalKill()*,  
*SyncCondvarSignal()*, *SyncMutexLock()*, *SyncTypeCreate()*,  
*ThreadCancel()*, *TimerTimeout()*

### Synopsis:

```
#include <sys/neutrino.h>

int SyncCtl(int cmd,
 sync_t * sync,
 void * data);


int SyncCtl_r(int cmd,
 sync_t * sync,
 void * data);
```

### Arguments:

*cmd* The operation type; one of:

- `_NTO_SCTL_GETPRIOCEILING` — get the ceiling priority of the mutex pointed to by *sync* and put it in the variable pointed to by *data*.
- `_NTO_SCTL_SETPRIOCEILING` — return the original ceiling priority. Set the ceiling priority of the mutex pointed to by *sync* to the value pointed to by *data*.
- `_NTO_SCTL_SETEVENT` — attach an event, pointed to by *data*, to the mutex pointed to by *sync*.

---

 You should use `SyncMutexEvent()` to do this.

---

*sync* A pointer to the synchronization object that you want to manipulate.

*data* A pointer to data associated with the command, or a place where the function can store the requested information, depending on the operation.

### Library:

`libc`

## Description:

The *SyncCtl()* and *SyncCtl\_r()* kernel calls let you:

- set or get a ceiling priority for a mutex  
or
- attach an event to a mutex so you'll be notified when the mutex changes to the DEAD state.

These functions are similar, except for the way they indicate errors. See the Returns section for details.

## Returns:

The only difference between these functions is the way they indicate errors:

|                    |                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>SyncCtl()</i>   | If an error occurs, the function returns -1 and sets <i>errno</i> . Any other value returned indicates success.                                           |
| <i>SyncCtl_r()</i> | Returns EOK on success. This function does <b>NOT</b> set <i>errno</i> . If an error occurs, the function returns any value listed in the Errors section. |

## Errors:

|        |                                                                                                                                              |
|--------|----------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | All kernel synchronization event objects are in use.                                                                                         |
| EFAULT | A fault occurred when the kernel tried to access <i>sync</i> or <i>data</i> .                                                                |
| EINVAL | The synchronization object pointed to by <i>sync</i> doesn't exist, or the ceiling priority value pointed to by <i>data</i> is out of range. |
| ENOSYS | The <i>SyncCtl()</i> and <i>SyncCtl_r()</i> functions aren't currently supported.                                                            |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*SyncCondvarSignal(), SyncCondvarWait(), SyncDestroy(),  
SyncMutexLock(), SyncMutexRevive(), SyncMutexUnlock(),  
SyncTypeCreate()*

# ***SyncDestroy()*, *SyncDestroy\_r()***

© 2004, QNX Software Systems Ltd.

*Destroy a synchronization object*

## **Synopsis:**

```
#include <sys/neutrino.h>

int SyncDestroy(sync_t* sync);

int SyncDestroy_r (sync_t* sync);
```

## **Arguments:**

*sync*     The synchronization object that you want to destroy.

## **Library:**

`libc`

## **Description:**

The *SyncDestroy()* and *SyncDestroy\_r()* kernel calls destroy a synchronization object previously allocated by a call to *SyncTypeCreate()*. If the object is a locked mutex, or a condition variable with waiting threads, the call fails. Any attempt to use *sync* after it is destroyed fails.

These functions are identical except in the way they indicate errors. See the Returns section for details.



Don't call *SyncDestroy()* directly; instead, use the POSIX synchronization objects (see *pthread\_cond\_destroy()*, *pthread\_mutex\_destroy()*, *pthread\_rwlock\_destroy()*, and *sem\_destroy()*).

---

## **Blocking states**

These calls don't block.



**Returns:**

The only difference between these functions is the way they indicate errors:

*SyncDestroy()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

*SyncDestroy\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function can return any value listed in the Errors section.

**Errors:**

- EBUSY      The synchronization object is locked by a thread.
- EFAULT     A fault occurred when the kernel tried to access *sync*.
- EINVAL     The synchronization ID specified in *sync* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## ***SyncDestroy(), SyncDestroy\_r()***

---

© 2004, QNX Software Systems Ltd.

### **See also:**

*SyncTypeCreate()*

## ***SyncMutexEvent()*, *SyncMutexEvent\_r()***

*Attach an event to a mutex*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SyncMutexEvent(sync_t * sync,
 struct sigevent * event);

int SyncMutexEvent_r(sync_t * sync,
 struct sigevent * event);
```

### **Arguments:**

*sync*      A pointer to the synchronization object for the mutex that you want to attach an event to.

*event*     A pointer to the **sigevent** structure that describes the event that you want to attach.

### **Library:**

**libc**

### **Description:**

The *SyncMutexEvent()* is a kernel call that attaches a specified *event* to a mutex pointed to by *sync*. You use *SyncMutexRevive()* to revive a DEAD mutex. Normally, a mutex will be placed in the DEAD state when the memory that was used to lock the mutex gets unmapped. One of the ways this may happen is when a process dies while holding the mutex in a shared memory.

*SyncMutexEvent()* and *SyncMutexEvent\_r()* are similar, except for the way they indicate errors. See the Returns section for details.

### **Returns:**

The only difference between these functions is the way they indicate errors:

## *SyncMutexEvent()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

## *SyncMutexEvent\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function returns any value listed in the Errors section.

## Errors:

|        |                                                                     |
|--------|---------------------------------------------------------------------|
| EAGAIN | All kernel synchronization event objects are in use.                |
| EFAULT | A fault occurred when the kernel tried to access <i>sync</i> .      |
| EINVAL | The synchronization object pointed to by <i>sync</i> doesn't exist. |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

**sigevent**, *SyncCondvarSignal()*, *SyncCondvarWait()*, *SyncDestroy()*, *SyncMutexLock()*, *SyncMutexRevive()*, *SyncMutexUnlock()*

## ***SyncMutexLock(), SyncMutexLock\_r()***

*Lock a mutex synchronization object*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SyncMutexLock(sync_t * sync);

int SyncMutexLock_r(sync_t * sync);
```

### **Arguments:**

*sync* A pointer to the synchronization object for the mutex that you want to lock.

### **Library:**

`libc`

### **Description:**

The *SyncMutexLock()* and *SyncMutexLock\_r()* kernel calls try to lock the mutex synchronization object *sync*. If the mutex isn't currently locked, the call returns immediately with the object locked. The mutex is considered unlocked if the owner field of *sync* is zero. Otherwise, the owner field of *sync* is treated as the thread ID of the current owner of the mutex.

These functions are similar, except for the way they indicate errors. See the Returns section for details.

If the mutex is already locked, the calling thread blocks on *sync* until it's unlocked by the owner. If more than one thread is blocked on *sync* they're queued in priority order.

If the priority of the blocking thread is higher than the thread that owns the mutex, the owner's priority is boosted to match that of the caller. In other words, the owner inherits the caller's priority if it's higher. If the owner's priority is boosted, it returns to its previous value before any boosts when the mutex is unlocked. Note that the owner may be boosted more than once as higher priority threads block on *sync*.

If a thread is boosted via this mechanism and subsequently changes its own priority, that priority takes immediate effect and also becomes the value it's returned to after it releases the mutex.

Waiting for a mutex isn't a cancellation point. If a signal is delivered to the thread while waiting for the mutex, the signal handler runs and, upon return from the handler, the thread resumes waiting for the mutex as if it wasn't interrupted.



---

Avoid timeouts on mutexes. Mutexes should be locked for brief periods of time, eliminating the need for a timeout.

---

The *sync* argument must have been initialized by a call to *SyncTypeCreate()* or have been statically initialized by the manifest `PTHREAD_MUTEX_INITIALIZER`.



---

The POSIX functions *pthread\_mutex\_lock()*, and *pthread\_mutex\_unlock()*, are faster, since they can potentially avoid a kernel call.

---

## Blocking states

`STATE_MUTEX`

The calling thread blocks waiting for the synchronization object to be unlocked.

## Returns:

The only difference between these functions is the way they indicate errors:

*SyncMutexLock()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

## SyncMutexLock\_r()

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, this function returns any value listed in the Errors section.

### Errors:

|           |                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------|
| EAGAIN    | On the first use of a statically initialized <i>sync</i> , all kernel synchronization objects were in use. |
| EFAULT    | A fault occurred when the kernel tried to access the buffers you provided.                                 |
| EINVAL    | The synchronization ID specified in <i>sync</i> doesn't exist.                                             |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                           |

### Classification:

QNX Neutrino

#### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_mutex\_lock()*, *pthread\_mutex\_unlock()*, *SyncTypeCreate()*, *SyncDestroy()*, *SyncMutexUnlock()*

# **SyncMutexRevive(), SyncMutexRevive\_r()**

© 2004, QNX

Software Systems Ltd.

*Revive a mutex that's in the DEAD state*

---

## **Synopsis:**

```
#include <sys/neutrino.h>

int SyncMutexRevive(sync_t * sync);

int SyncMutexRevive_r(sync_t * sync);
```

## **Arguments:**

*sync*     A pointer to the synchronization object for the mutex that you want to revive.

## **Library:**

`libc`

## **Description:**

The *SyncMutexRevive()* and *SyncMutexRevive\_r()* kernel calls revive a mutex, pointed to by *sync*, that's in the DEAD state. The mutex will be put into the LOCK state and will be owned by the calling thread. The mutex counter is set to one (for recursive mutexes).

These functions are similar, except for the way they indicate errors. See the Returns section for details.

See *SyncMutexEvent()* for information on how to get notified when a mutex enters the DEAD state.

## **Returns:**

The only difference between these functions is the way they indicate errors:

*SyncMutexRevive()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.



## *SyncMutexRevive\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function returns any value listed in the Errors section.

### Errors:

|           |                                                                                                 |
|-----------|-------------------------------------------------------------------------------------------------|
| EFAULT    | A fault occurred when the kernel tried to access the buffers you provided.                      |
| EINVAL    | The synchronization object pointed to by <i>sync</i> doesn't exist or wasn't in the DEAD state. |
| ETIMEDOUT | A kernel timeout unblocked the call. See <i>TimerTimeout()</i> .                                |

### Classification:

QNX Neutrino

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*pthread\_mutex\_lock()*, *pthread\_mutex\_unlock()*, *SyncTypeCreate()*, *SyncDestroy()*, *SyncMutexEvent()*, *SyncMutexLock()*, *SyncMutexUnlock()*

# SyncMutexUnlock(), SyncMutexUnlock\_r()

© 2004, QNX

Software Systems Ltd.

Unlock a mutex synchronization object

## Synopsis:

```
#include <sys/neutrino.h>

int SyncMutexUnlock(sync_t * sync);

int SyncMutexUnlock_r(sync_t * sync);
```

## Arguments:

*sync* A pointer to the synchronization object for the mutex that you want to unlock.

## Library:

libc

## Description:

The *SyncMutexUnlock()* and *SyncMutexUnlock\_r()* kernel calls unlock the mutex passed as *sync*. If there are threads blocked on the mutex, the *owner* member of *sync* is set to the thread ID of the thread with the highest priority that has been waiting the longest and it's made ready to run. If no threads are waiting, it's set to zero.

These functions are similar, except for the way they indicate errors. See the Returns section for details.

If the calling thread had its priority boosted while it owned the mutex, it returns to its normal priority.



---

The POSIX functions *pthread\_mutex\_lock()*, and *pthread\_mutex\_unlock()*, are faster, since they can potentially avoid a kernel call.

---

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *SyncMutexUnlock()*

If an error occurs, The function returns -1 and sets *errno*. Any other value returned indicates success.

### *SyncMutexUnlock\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function returns any value listed in the Errors section.

## Errors:

- |        |                                                                                                          |
|--------|----------------------------------------------------------------------------------------------------------|
| EFAULT | A fault occurred when the kernel tried to access the buffers provided.                                   |
| EINVAL | The synchronization ID specified in <i>sync</i> doesn't exist. The calling thread doesn't own the mutex. |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## ***SyncMutexUnlock()*, *SyncMutexUnlock\_r()***

© 2004, QNX

Software Systems Ltd.

---

### **See also:**

*pthread\_mutex\_lock()*, *pthread\_mutex\_unlock()*, *SyncTypeCreate()*,  
*SyncMutexLock()*

## ***SyncSemPost()*, *SyncSemPost\_r()***

*Increment a semaphore*

### **Synopsis:**

```
#include <sys/neutrino.h>

int SyncSemPost(sync_t* sync);

int SyncSemPost_r(sync_t* sync);
```

### **Arguments:**

*sync*     A pointer to the synchronization object for the semaphore that you want to increment.

### **Library:**

`libc`

### **Description:**

The *SyncSemPost()* and *SyncSemPost\_r()* kernel calls increment the semaphore referenced by the *sync* argument. If any threads are blocked on the semaphore, the one waiting the longest is unblocked and allowed to run.

These functions are identical, except for the way they indicate errors. See the Returns section for details.



You should use the POSIX *sem\_post()* function instead of calling *SyncSemPost()* directly.

---

### **Returns:**

The only difference between these functions is the way they indicate errors:

*SyncSemPost()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

*SyncSemPost\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function returns one of the values listed in the Errors section.

**Errors:**

|        |                                                                             |
|--------|-----------------------------------------------------------------------------|
| EAGAIN | Not enough memory for the kernel to create the internal <i>sync</i> object. |
| EFAULT | Invalid pointer.                                                            |
| EINTR  | A signal interrupted this function.                                         |
| EINVAL | The <i>sync</i> argument doesn't refer to a valid semaphore.                |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sem\_destroy()*, *sem\_init()*, *sem\_post()*, *sem\_trywait()*, *sem\_wait()*, *SyncDestroy()*, *SyncSemWait()*, *SyncTypeCreate()*

**Synopsis:**

```
#include <sys/neutrino.h>

int SyncSemWait(sync_t* sync,
 int try);

int SyncSemWait_r(sync_t* sync,
 int try);
```

**Arguments:**

*sync*     A pointer to the synchronization object for the semaphore that you want to wait on.

*try*     Nonzero if you want a conditional wait.

**Library:**

`libc`

**Description:**

The *SyncSemWait()* and *SyncSemWait\_r()* kernel calls decrement the semaphore referred to by the *sync* argument. If the semaphore value isn't greater than zero and *try* is zero, then the calling process blocks until it can decrement the counter or the call is interrupted by signal.

These functions are identical, except in the way they indicate errors. See the Returns section for details.

If *try* is nonzero, the function acts as a conditional wait. If the call would block, the semaphore is unmodified, and the call returns with an error.

**Returns:**

The only difference between these functions is the way they indicate errors:

*SyncSemWait()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success (the semaphore was successfully decremented).

*SyncSemWait\_r()*

Returns EOK on success (the semaphore was successfully decremented). This function does **NOT** set *errno*. If an error occurs, the function returns one of the values listed in the Errors section.

**Errors:**

|         |                                                              |
|---------|--------------------------------------------------------------|
| EAGAIN  | Call would have blocked and <i>try</i> was nonzero.          |
| EDEADLK | A deadlock condition was detected.                           |
| EINTR   | A signal interrupted this function.                          |
| EINVAL  | The <i>sync</i> argument doesn't refer to a valid semaphore. |

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*sem\_destroy(), sem\_init(), sem\_post(), sem\_trywait(), sem\_wait(),  
SyncDestroy(), SyncSemPost(), SyncTypeCreate()*

# SyncTypeCreate(), SyncTypeCreate\_r()

© 2004, QNX Software

Systems Ltd.

Create a synchronization object

## Synopsis:

```
#include <sys/neutrino.h>

int SyncTypeCreate(
 unsigned type,
 sync_t * sync,
 const struct _sync_attr_t * attr);

int SyncTypeCreate_r(
 unsigned type,
 sync_t * sync,
 const struct _sync_attr_t * attr);
```

## Arguments:

- type* One of the following:
- `_NTO_SYNC_MUTEX_FREE` — create a mutex.
  - `_NTO_SYNC_SEM` — create a semaphore.
  - `_NTO_SYNC_COND` — create a condition variable.
- sync* A pointer to a `sync_t` that the kernel sets up for the synchronization object; see below.
- attr* A pointer to a `_sync_attr_t` structure that specifies attributes for the object. This structure contains at least the following members:
- `int protocol` — `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT`.
- If *attr* is `NULL`, the default attributes (`PTHREAD_PRIO_INHERIT`) are assumed.

## Library:

`libc`

**Description:**

The *SyncTypeCreate()* and *SyncTypeCreate\_r()* kernel calls create a synchronization object in the kernel and initializes *sync* for use in other synchronization kernel calls. The synchronization object is local to the process.

These functions are similar, except for the way they indicate errors. See the Returns section for details.

Synchronization objects can be used for mutexes, semaphores, or condition variables.




---

Don't call *SyncTypeCreate()* directly; instead, use the POSIX synchronization objects (see *pthread\_cond\_init()*, *pthread\_mutex\_init()*, *pthread\_rwlock\_init()*, and *sem\_init()*).

---

The *sync* argument contains at least the following members:

- |                         |                                                                                                                                                                                            |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>int</b> <i>count</i> | The count for recursive mutexes and semaphores. The kernel sets this member when it creates the synchronization object.                                                                    |
| <b>int</b> <i>owner</i> | When a mutex is created, this member holds the thread ID of the thread that acquired the mutex. When unowned, the value is 0. It's set to zero when the synchronization object is created. |

The current state of *sync* is summarized below:

| <b>Counter</b> | <b>Owner</b> | <b>Description</b>                             |
|----------------|--------------|------------------------------------------------|
| -              | -2           | Destroyed mutex                                |
| 0              | -1           | Statically initialized; auto-created when used |

*continued...*

| Counter      | Owner | Description                                                                  |
|--------------|-------|------------------------------------------------------------------------------|
| 0            | 0     | Unlocked mutex                                                               |
| <i>count</i> | >0    | Recursive counter number of the mutex                                        |
| <i>count</i> | <-1   | If the high bit of <i>count</i> is set, it's a flag meaning "others waiting" |
| -            | -256  | Mutex is dead, waits for revival                                             |

The synchronization object is destroyed by a call to *SyncDestroy()*.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *SyncTypeCreate()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

### *SyncTypeCreate\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function can return any value in the Errors section.

## Errors:

|        |                                                                               |
|--------|-------------------------------------------------------------------------------|
| EAGAIN | All kernel synchronization objects are in use.                                |
| EFAULT | A fault occurred when the kernel tried to access <i>sync</i> or <i>attr</i> . |
| EINVAL | Either                                                                        |

- the *type* isn't one of `_NTO_SYNC_COND`, `_NTO_SYNC_MUTEX_FREE` or `_NTO_SYNC_SEM`  
Or:
- if the type is correct, and the synchronization object is:
  - a mutex — the protocol isn't one of `PTHREAD_PRIO_INHERIT` or `PTHREAD_PRIO_PROTECT`.
  - a mutex and `PTHREAD_PRIO_PROTECT` is specified — the ceiling priority isn't within the kernel priority range.
  - a condvar — the clock type is invalid.
  - a semaphore — the semaphore value exceeds `SEM_VALUE_MAX`.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*pthread\_cond\_init(), pthread\_mutex\_init(), pthread\_rwlock\_init(), sem\_init(), SyncCondvarSignal(), SyncCondvarWait(), SyncDestroy(), SyncMutexLock(), SyncMutexUnlock()*

## **sysconf()**

© 2004, QNX Software Systems Ltd.

*Return the value of a configurable system limit*

### **Synopsis:**

```
#include <unistd.h>
#include <limits.h>

long sysconf(int name);
```

### **Arguments:**

*name*     The name of the limit that you want to get; see below.

### **Library:**

`libc`

### **Description:**

The *sysconf()* function returns the value of a configurable system limit indicated by *name*.

Configurable limits are defined in `<confname.h>`, and contain at least the following values:

`_SC_ARG_MAX`

Maximum length of arguments for the *exec\*()* functions, in bytes, including environment data.

`_SC_CHILD_MAX`

Maximum number of simultaneous processes per real user ID.

`_SC_CLK_TCK`

The number of intervals per second used to express the value in type `clock_t`.

`_SC_NGROUPS_MAX`

The maximum number of simultaneous supplementary group IDs per process.

**\_SC\_OPEN\_MAX**

Maximum number of files that one process can have open at any given time.

**\_SC\_JOB\_CONTROL**

If this variable is defined, then job control is supported.

**\_SC\_SAVED\_IDS**

If this variable is defined, then each process has a saved set-user ID and a saved set-group ID.

**\_SC\_VERSION**

The current POSIX version that is currently supported. A value of 198808L indicates the August (08) 1988 standard, as approved by the IEEE Standards Board.

**Returns:**

The requested configurable system limit. If *name* isn't defined for the system, -1 is returned.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <limits.h>
#include <unistd.h>

int main(void)
{
 printf("_SC_ARG_MAX = %ld\n",
 sysconf(_SC_ARG_MAX));
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*confstr(), errno, pathconf()*

**getconf** in the *Utilities Reference*

Understanding System Limits chapter of the Neutrino *User's Guide*



**Synopsis:**

```
#include <sys/param.h>
#include <sys/sysctl.h>

int sysctl(int * name,
 u_int namelen,
 void * oldp,
 size_t * oldlenp,
 void * newp,
 size_t newlen);
```

**Arguments:**

|                |                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>name</i>    | An array of integers that specifies the <i>Management Information Base</i> (MIB) stylename of the item that you want to set or get; see below.                                                                       |
| <i>namelen</i> | The length of the name.                                                                                                                                                                                              |
| <i>oldp</i>    | NULL, or a pointer to a buffer where the function can store the old value.                                                                                                                                           |
| <i>oldlenp</i> | NULL, or a pointer to a location that initially specifies the size of the <i>oldp</i> buffer. The function changes the value in this location to be the size of the old information stored in the <i>oldp</i> buffer |
| <i>newp</i>    | NULL, or a pointer to a buffer that holds the new value.                                                                                                                                                             |
| <i>newlen</i>  | The size of the new value.                                                                                                                                                                                           |

**Library:**

**libsocket**

## Description:

The *sysctl()* function retrieves system information and allows processes with appropriate privileges to set system information. The data available from *sysctl()* consists of integers and tables. You can also get or set data using the **sysctl** utility at the command line.

The state is described using a *Management Information Base* (MIB) stylename, specified in *name*, which is a *namelen* length array of integers.

The information is copied into the buffer specified by *oldp*. The size of the buffer is given by the location specified by *oldlenp* before the call, and that location gives the amount of data copied after a successful call. If the amount of data available is greater than the size of the buffer supplied, the call delivers as much data as fits in the buffer provided and returns with the error code ENOMEM. If you don't need the old value, you can set *oldp* and *oldlenp* to NULL.

You can determine the size of the available data by calling *sysctl()* with a NULL parameter for *oldp*. The function stores the size of the available data in the location pointed to by *oldlenp*. For some operations, the amount of space may change often. For these operations, the system attempts to round up, so that the returned size is large enough for a call to return the data shortly thereafter.

To specify a new value, set *newp* to point to a buffer of length *newlen* from which the requested value is to be taken. If you're not setting a new value, set *newp* to NULL and *newlen* to 0.

The top-level names are defined with a CTL\_ prefix in `<sys/sysctl.h>`. QNX 4 supports CTL\_NET only. The next and subsequent levels down are found in the following header files:

| This header file                  | Contains definitions for         |
|-----------------------------------|----------------------------------|
| <code>&lt;sys/sysctl.h&gt;</code> | Top-level identifiers            |
| <code>&lt;sys/socket.h&gt;</code> | Second-level network identifiers |

*continued...*

| This header file     | Contains definitions for                                         |
|----------------------|------------------------------------------------------------------|
| <netinet/in.h>       | Third-level Internet identifiers and fourth-level IP identifiers |
| <netinet/icmp_var.h> | Fourth-level ICMP identifiers                                    |
| <netinet/tcp_var.h>  | Fourth-level TCP identifiers                                     |
| <netinet/udp_var.h>  | Fourth-level UDP identifiers                                     |

The following code fragment checks whether the UDP packets checksum is enabled:

```
int mib[5], val;
size_t len;

mib[0] = CTL_NET;
mib[1] = AF_INET;
mib[2] = IPPROTO_UDP;
mib[3] = UDPCTL_CHECKSUM;
len = sizeof(val);
sysctl(mib, 4, &val, &len, NULL, 0);
```

## CTL\_NET

The table and integer information available for the CTL\_NET level is detailed below. The changeable column shows whether a process with appropriate privilege may change the value.

| Second-level name | Type            | Changeable |
|-------------------|-----------------|------------|
| PF_INET           | internet values | yes        |

## PF\_INET

PF\_INET gets or sets global information about internet protocols.

The third-level name is the protocol. The fourth-level name is the variable name. Here are the currently defined protocols and names:

| Protocol name | Variable name             | Type    | Changeable |
|---------------|---------------------------|---------|------------|
| <i>ip</i>     | <i>forwarding</i>         | integer | yes        |
|               | <i>redirect</i>           | integer | yes        |
|               | <i>ttl</i>                | integer | yes        |
|               | <i>forwsrct</i>           | integer | yes        |
|               | <i>directed-broadcast</i> | integer | yes        |
|               | <i>allowsrct</i>          | integer | yes        |
|               | <i>subnetsarelocal</i>    | integer | yes        |
|               | <i>mtudisc</i>            | integer | yes        |
|               | <i>maxfragpackets</i>     | integer | yes        |
|               | <i>sourcecheck</i>        | integer | yes        |
|               | <i>sourcecheck_logint</i> | integer | yes        |
| <i>icmp</i>   | <i>maskrepl</i>           | integer | yes        |
| <i>tcp</i>    | <i>rfc1323</i>            | integer | yes        |
|               | <i>sendspace</i>          | integer | yes        |
|               | <i>recvspace</i>          | integer | yes        |
|               | <i>mssdflt</i>            | integer | yes        |
|               | <i>syn_cache_limit</i>    | integer | yes        |
|               | <i>syn_bucket_limit</i>   | integer | yes        |
| <i>udp</i>    | <i>checksum</i>           | integer | yes        |
|               | <i>sendspace</i>          | integer | yes        |
|               | <i>recvspace</i>          | integer | yes        |

The variables are as follows:

|                              |                                                                                                                                                                              |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ip.forwarding</i>         | Returns 1 when IP forwarding is enabled for the host, meaning that the host is acting as a router.                                                                           |
| <i>ip.redirect</i>           | Returns 1 when ICMP redirects may be sent by the host. This option is ignored unless the host is routing IP packets. Normally, this option should be enabled on all systems. |
| <i>ip.ttl</i>                | The maximum time-to-live (hop count) value for an IP packet sourced by the system. This value applies to normal transport protocols, not to ICMP.                            |
| <i>ip.forwsrct</i>           | Returns 1 when forwarding of source-routed packets is enabled for the host. This value may be changed only if the kernel security level is less than 1.                      |
| <i>ip.directed-broadcast</i> | Returns 1 if directed-broadcast behavior is enabled for the host.                                                                                                            |
| <i>ip.allowsrct</i>          | Returns 1 if the host accepts source-routed packets.                                                                                                                         |
| <i>ip.subnetsarelocal</i>    | Returns 1 if subnets are to be considered local addresses.                                                                                                                   |
| <i>ip.mtudisc</i>            | Returns 1 if path MTU discovery is enabled.                                                                                                                                  |
| <i>ip.maxfragpackets</i>     | Returns the maximum number of fragmented IP packets in the IP reassembly queue.                                                                                              |
| <i>ip.sourcecheck</i>        | Returns 1 if source check for received packets is enabled.                                                                                                                   |
| <i>ip.sourcecheck_logint</i> | Returns the time interval when IP source address verification messages are logged. A value of zero disables the logging.                                                     |

|                               |                                                                                                   |
|-------------------------------|---------------------------------------------------------------------------------------------------|
| <i>icmp.maskrepl</i>          | Returns 1 if ICMP network mask requests are to be answered.                                       |
| <i>tcp.rfc1323</i>            | Returns 1 if RFC1323 extensions to TCP are enabled.                                               |
| <i>tcp.sendspace</i>          | Returns the default TCP send buffer size.                                                         |
| <i>tcp.recvspace</i>          | Returns the default TCP receive buffer size.                                                      |
| <i>tcp.mssdflt</i>            | Returns the default TCP maximum segment size.                                                     |
| <i>tcp.syn_cache_limit</i>    | Returns the maximum number of entries allowed in the TCP compressed state engine.                 |
| <i>tcp.syn_bucket_limit</i>   | Returns the maximum number of entries allowed per hash bucket in the TCP compressed state engine. |
| <i>tcp.syn_cache_interval</i> | Returns the TCP compressed state engine's timer interval.                                         |
| <i>udp.checksum</i>           | Returns 1 when UDP checksums are being computed and checked.                                      |



---

Disabling UDP checksums is strongly discouraged.

---

|                      |                                              |
|----------------------|----------------------------------------------|
| <i>udp.sendspace</i> | Returns the default UDP send buffer size.    |
| <i>udp.recvspace</i> | Returns the default UDP receive buffer size. |

## Returns:

|    |                                           |
|----|-------------------------------------------|
| 0  | Success.                                  |
| -1 | An error occurred ( <i>errno</i> is set). |

## Errors:

|            |                                                                                                                                                                           |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EFAULT     | The buffers: <i>name</i> , <i>oldp</i> , <i>newp</i> , or the length pointer <i>oldlenp</i> contains an invalid address.                                                  |
| EINVAL     | The name array is less than two or greater than CTL_MAXNAME; or a non-NULL <i>newp</i> is given and its specified length in <i>newlen</i> is too large or too small.      |
| ENOMEM     | The length pointed to by <i>oldlenp</i> is too short to hold the requested value.                                                                                         |
| ENOTDIR    | The name array specifies an intermediate rather than terminal name.                                                                                                       |
| EOPNOTSUPP | The name array specifies an unknown value.                                                                                                                                |
| EPERM      | An attempt was made to set a read-only value; a process, without appropriate privilege, attempts to set or change a value protected by the current system security level. |

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

ROUTE protocol

**sysctl** in the *Utilities Reference*



## Synopsis:

```
#include <syslog.h>

void syslog(int priority,
 const char * message,
 ...)
```

## Arguments:

- priority*      The priority of the message; see “Message levels,” below.
- message*      The message that you want to write. This message is identical to a *printf()*-format string, except that *%m* is replaced by the current error message (as denoted by the global variable *errno*). A trailing newline is added if none is present.

The formatting characters that you use in the message determine any additional arguments.

## Library:

`libc`

## Description:

The *syslog()* function writes *message* to the system message logger. The message is then written to the system console, log files, and logged-in users, or forwarded to other machines as appropriate. (See the `syslogd` command.)

The *vsyslog()* function is an alternate form in which the arguments have already been captured using the variable-length argument facilities of `<stdarg.h>`.

## Message levels

The message is tagged with *priority*. Priorities are encoded as a *facility* and a *level*. The facility describes the part of the system generating the message. The level is selected from the following list (ordered from high to low):

|             |                                                                                        |
|-------------|----------------------------------------------------------------------------------------|
| LOG_EMERG   | A panic condition. This is normally broadcast to all users.                            |
| LOG_ALERT   | A condition that should be corrected immediately, such as a corrupted system database. |
| LOG_CRIT    | Critical conditions (for example, hard device errors).                                 |
| LOG_ERR     | General errors.                                                                        |
| LOG_WARNING | Warning messages.                                                                      |
| LOG_NOTICE  | Conditions that aren't error conditions, but should possibly be specially handled.     |
| LOG_INFO    | Informational messages.                                                                |
| LOG_DEBUG   | Messages that contain information normally of use only when debugging a program.       |

## Examples:

```
syslog(LOG_ALERT, "who: internal error 23");

openlog("ftpd", LOG_PID, LOG_DAEMON);
setlogmask(LOG_UPTO(LOG_ERR));
syslog(LOG_INFO, "Connection from host %d", CallingHost);

syslog(LOG_INFO|LOG_LOCAL2, "foobar error: %m");
```

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## See also:

*closelog(), openlog(), setlogmask(), vsyslog()*

**logger**, **syslogd** in the *Utilities Reference*

# ***sysmgr\_reboot()***

© 2004, QNX Software Systems Ltd.

*Reboot the system*

## **Synopsis:**

```
#include <sys/sysmgr.h>

int sysmgr_reboot(void);
```

## **Library:**

libc

## **Description:**

The *sysmgr\_reboot()* function reboots the calling computer. You need to be **root** for this function to succeed.

## **Returns:**

The *sysmgr\_reboot()* function doesn't return if successful. If an error occurs, it returns:

EPERM      You need to be **root** to call this function.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`procnto` in the *Utilities Reference*

# ***SYSPAGE\_CPU\_ENTRY()***

© 2004, QNX Software Systems Ltd.

*Return a CPU-specific entry from the system page*

## **Synopsis:**

```
#include <sys/syspage.h>

#define SYSPAGE_CPU_ENTRY(cpu, entry)...
```

## **Arguments:**

*cpu*      The CPU to get the entry for.

*entry*     The entry to get; see below.

## **Library:**

`libc`

## **Description:**

The *SYSPAGE\_CPU\_ENTRY()* macro returns a pointer to the specified *entry* from the part of the system page that's specific to the given *cpu*.

The best way to reference the system page is via the kernel calls and POSIX cover functions. If there isn't a function to access the information you need, use *SYSPAGE\_CPU\_ENTRY()* instead of referencing the *\_syspage\_ptr* variable directly. For information in the rest of the *syspage\_entry* structure, use *SYSPAGE\_ENTRY()*.

The only entry you might currently need to use is:

### **ppc, kerinfo**

This structure, defined in `<ppc/syspage.h>`, contains at least the following members:

- **unsigned long** *pretend\_cpu* — we can pretend the chip is this Processor Version Register.
- **unsigned long** *init\_msr* — the initial Machine Status Register for thread creation.

**Returns:**

A pointer to the structure for the given *entry*.

**Examples:**

```
#include <inttypes.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/syspage.h>

int main(void)
{
 printf ("We're pretending to be a type %ld PPC\n",
 SYSPAGE_CPU_ENTRY (ppc, kerinfo) ->pretend_cpu);

 return EXIT_SUCCESS;
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*SYSPAGE\_CPU\_ENTRY()* is a macro.

**See also:**

*ClockCycles()*, *SYSPAGE\_ENTRY()*, *\_syspage\_ptr*

Customizing Image Startup Programs chapter of the *Building Embedded Systems* guide

# ***SYSPAGE\_ENTRY()***

© 2004, QNX Software Systems Ltd.

*Return an entry from the system page*

## **Synopsis:**

```
#include <sys/syspage.h>

#define SYSPAGE_ENTRY(entry)...
```

## **Arguments:**

*entry*     The entry to get; see below.

## **Library:**

`libc`

## **Description:**

The *SYSPAGE\_ENTRY()* macro returns a pointer to the specified *entry* in the system page.

The best way to reference the system page is via the kernel calls and POSIX cover functions. If there isn't a function to access the information you need, use *SYSPAGE\_ENTRY()* instead of referencing the *\_syspage\_ptr* variable directly. For information in the CPU-specific part of the `syspage_entry` structure, use *SYSPAGE\_CPU\_ENTRY()*.

Currently, the only *entry* you're likely to access with *SYSPAGE\_ENTRY()* is:

- qtime**     QNX-specific time information. The `qtime_entry` structure contains at least the following members:
- **unsigned long** *boot\_time* — the time, in seconds, since the Unix Epoch (00:00:00 January 1, 1970 Coordinated Universal Time (UTC)) when this system was booted.
  - **uint64\_t** *cycles\_per\_sec* — the number of CPU clock cycles per second for this system. For more information, see *ClockCycles()*.



**Returns:**

A pointer to the structure for the given *entry*.

**Examples:**

```
#include <sys/neutrino.h>
#include <inttypes.h>
#include <stdio.h>
#include <stdlib.h>
#include <sys/syspage.h>

int main(void)
{
 uint64_t cps, cycle1, cycle2, ncycles;
 float sec;

 /* snap the time */
 cycle1=ClockCycles();

 /* do something */
 printf("poo\n");

 /* snap the time again */
 cycle2=ClockCycles();
 ncycles=cycle2-cycle1;
 printf("%lld cycles elapsed \n", ncycles);

 /* find out how many cycles per second */
 cps = SYSPAGE_ENTRY(qtime)->cycles_per_sec;
 printf("This system has %lld cycles/sec.\n",cps);
 sec=(float)ncycles/cps;
 printf("The cycles in seconds is %f \n",sec);

 return EXIT_SUCCESS;
}
```

**Classification:**

QNX Neutrino

**Safety**

Cancellation point No

*continued...*

## **Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

## **Caveats:**

*SYSPAGE\_ENTRY()* is a macro.

## **See also:**

*ClockCycles()*, *SYSPAGE\_CPU\_ENTRY()*, *\_syspage\_ptr*

Customizing Image Startup Programs chapter of the *Building Embedded Systems* guide

## Synopsis:

```
#include <sys/syspage.h>

struct syspage_entry *_syspage_ptr
```

## Description:

This global variable holds a pointer to the *system page*, which contains information about the system, including the processor type, bus type, and the location and size of available system RAM. For information about this structure, see the Customizing Image Startup Programs chapter of *Building Embedded Systems*.

The best way to reference the system information page is via the kernel calls and POSIX cover functions. If there isn't a function to access the information you need, you should use the *SYSPAGE\_ENTRY()* and *SYSPAGE\_CPU\_ENTRY()* macros instead of referencing the *\_syspage\_ptr* variable directly.

## Classification:

QNX Neutrino

## See also:

*SYSPAGE\_CPU\_ENTRY()*, *SYSPAGE\_ENTRY()*

## **system()**

© 2004, QNX Software Systems Ltd.

*Execute a system command*

---

### **Synopsis:**

```
#include <stdlib.h>

int system(const char *command);
```

### **Arguments:**

*command*      NULL, or the system command that you want to execute; see below.

### **Library:**

`libc`

### **Description:**

The behavior of the *system()* function depends on the value of its *command* argument:

- If *command* is NULL, *system()* determines whether or not a shell is present.
- If *command* isn't NULL, *system()* invokes a copy of the shell, and passes the string *command* to it for processing. This function uses *spawnlp()* to load a copy of the shell.



The shell used is always `/bin/sh`, regardless of the setting of the **SHELL** environment variable, because applications may rely on features of the standard shell, and may fail as a result of running a different shell.

This means that any command that can be entered to the OS can be executed, including programs, QNX Neutrino commands, and shell scripts. The *exec\*()* and *spawn\*()* functions can only cause programs to be executed.

---

## Returns:

- If *command* is NULL, *system()* returns zero if the shell isn't present, or a nonzero value if the shell is present.
- If *command* isn't NULL, *system()* returns the result of invoking a copy of the shell. If the shell couldn't be loaded, *system()* returns -1; otherwise, it returns the status of the specified command. Use the *WEXITSTATUS()* macro to determine the low-order 8 bits of the termination status of the process.

For example, assume that *status* is the value returned by *system()*. If *WEXITSTATUS( status )* returns 255, either the specified command returned a termination status of 255, or the shell didn't exit (i.e. it died from a signal or couldn't be started at all) and the return value was 255 due to implementation details. For example, under QNX Neutrino and most Unix systems, the value is 255 if *status* is -1, which indicates that the shell couldn't be executed. *WEXITSTATUS()* is defined in `<sys/wait.h>`.

For information about macros that extract information from the value returned by *system()*, see "Status macros" in the description of *wait()*.

When an error has occurred, *errno* contains a value that indicates the type of error that has been detected.

## Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <sys/wait.h>

int main(void)
{
 int rc;

 rc = system("ls");
 if(rc == -1) {
 printf("shell could not be run\n");
 } else {
 printf("result of running command is %d\n",
 WEXITSTATUS(rc));
 }
 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*abort(), atexit(), close(), errno, execl(), execlp(), execlpe(),  
execv(), execve(), execvp(), execvpe(), exit(), \_exit(), getenv(), main(),  
putenv(), sigaction(), signal(), spawn(), spawnl(), spawnle(),  
spawnlp(), spawnlpe(), spawnp(), spawnv(), spawnve(), spawnvp(),  
spawnvpe(), wait(), waitpid()*

## Synopsis:

```
#include <math.h>

double tan(double x);

float tanf(float x);
```

## Arguments:

*x* The angle, in radians, for which you want to compute the tangent.

## Library:

libm

## Description:

These functions compute the tangent (specified in radians) of *x*. A large magnitude argument may yield a result with little or no significance.

## Returns:

The tangent value. If *x* is NAN or infinity, NAN is returned.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
{
 printf("%f\n", tan(.5));
}
```

```
 return EXIT_SUCCESS;
}
```

produces the output:

0.546302

### Classification:

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*atan()*, *atan2()*, *cos()*, *sin()*



## Synopsis:

```
#include <math.h>

double tanh(double x);

float tanhf(float x);
```

## Arguments:

*x* The angle, in radians, for which you want to compute the hyperbolic tangent.

## Library:

**libm**

## Description:

These functions compute the hyperbolic tangent (specified in radians) of *x*.

When the *x* argument is large, partial or total loss of significance may occur.

## Returns:

The hyperbolic tangent value.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

int main(void)
```

```
{
 printf("%f\n", tanh(.5));
 return EXIT_SUCCESS;
}
```

produces the output:

0.462117

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### See also:

*cosh()*, *errno*, *sinh()*

## Synopsis:

```
#include <termios.h>

int tcdrain(int fdes);
```

## Arguments:

*fdes*     A file descriptor that's associated with the device that you want to wait for.

## Library:

libc

## Description:

The *tcdrain()* function waits until all output has been physically transmitted to the device associated with *fdes*, or until a signal is received.

## Returns:

0     Success.

-1     An error occurred (*errno* is set).

## Errors:

EBADF     The argument *fdes* is invalid.

EINTR     A signal interrupted the operation.

ENOSYS     The resource manager associated with *fdes* doesn't support this call.

ENOTTY     The argument *fdes* doesn't refer to a terminal device.

**Examples:**

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 int fildes;

 fildes = open("/dev/ser1", O_RDWR);
 write(fildes, "ATH", 3);

 /* Wait for data to transmit before returning */
 tcdrain(fildes);
 close(fildes);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*write()*

## Synopsis:

```
#include <termios.h>

int tcdropline(int fd,
 int duration);
```

## Arguments:

*fd*            A file descriptor that's associated with the line that you want to disconnect.

*duration*     The number of milliseconds that you want to drop the line for.

## Library:

`libc`

## Description:

The *tcdropline()* function initiates a disconnect condition on the communication line associated with the opened file descriptor indicated by *fd*.

The disconnect condition lasts at least *duration* milliseconds, or approximately 300 milliseconds if *duration* is zero. The system rounds the effective value of *duration* up to the next highest supported interval, which is typically a multiple of 100 milliseconds.

## Returns:

0            Success.

-1          An error occurred (*errno* is set).

**Errors:**

- EBADF      Invalid *fd* argument.
- ENOSYS     The resource manager associated with *fd* doesn't support this call.
- ENOTTY     The argument *fd* doesn't refer to a terminal device.

**Examples:**

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 int fd;

 fd = open("/dev/ser1", O_RDWR);

 /* Disconnect for 500 milliseconds */
 tcdropline(fd, 500);

 close(fd);
 return EXIT_SUCCESS;
}
```

**Classification:**

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*tcdrain(), tcflow(), tcflush()*

## ***tcf*low()**

© 2004, QNX Software Systems Ltd.

*Perform a flow-control operation on a data stream*

### **Synopsis:**

```
#include <termios.h>

int tcf
```

*low*

```
(int fildes,
 int action);
```

### **Arguments:**

*fildes*     A file descriptor that's associated with the data stream that you want to perform the operation on.

*action*     The action you want to perform; see below.

### **Library:**

`libc`

### **Description:**

The *tcf*low() function performs a flow-control operation on the data stream associated with *fildes*, depending on the values in *action*.

At least the following actions are defined in `<termios.h>`:

|          |                                                                                                                                                                           |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TCOOFF   | Use software flow control to suspend output on the device associated with <i>fildes</i> .                                                                                 |
| TCOOFFHW | Use hardware flow control to suspend output on the device associated with <i>fildes</i> .                                                                                 |
| TCOON    | Use software flow control to resume output on the device associated with <i>fildes</i> .                                                                                  |
| TCOONHW  | Use hardware flow control to resume output on the device associated with <i>fildes</i> .                                                                                  |
| TCIOFF   | Cause input to be flow-controlled by sending a STOP character immediately across the communication line associated with <i>fildes</i> , (that is, software flow control). |



|          |                                                                                                                                                     |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| TCIOFFHW | Cause input to be flow-controlled by using hardware control.                                                                                        |
| TCION    | Resume input by sending a START character immediately across the communication line associated with <i>fildev</i> (that is, software flow control). |
| TCIONHW  | Cause input to be resumed by using hardware flow control.                                                                                           |

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EBADF Invalid *fildev* argument.
- EINVAL Invalid *action* argument.
- ENOSYS The resource manager associated with *fildev* doesn't support this call.
- ENOTTY The argument *fildev* doesn't refer to a terminal device.

## Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 int fd;

 fd = open("/dev/ser1", O_RDWR);

 /* Resume output on flow-controlled device */
 tcflow(fd, TCOON);
}
```

```
 close(fd);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*tcdrain()*, *tcflush()*, *tcsendbreak()*

## Synopsis:

```
#include <termios.h>

int tcflush(int fdes,
 int queue_selector);
```

## Arguments:

*fdes*      A file descriptor that's associated with the data stream that you want to perform the operation on.

*queue\_selector*

The stream or streams that you want to flush. At least the following values for *queue\_selector* are defined in `<termios.h>`:

- TCIFLUSH — discard all data that's received, but not yet read, on the device associated with *fdes*.
- TCOFLUSH — discard all data that's written, but not yet transmitted, on the device associated with *fdes*.
- TCIOFLUSH — discard all data that's written, but not yet transmitted, as well as all data that's received, but not yet read, on the device associated with *fdes*.

## Library:

`libc`

## Description:

The *tcflush()* function flushes the input stream, the output stream, or both, depending on the value of the argument *queue\_selector*.

## Returns:

- 0      Success.
- 1     An error occurred (*errno* is set).

**Errors:**

- EBADF      Invalid *fildev* argument.
- EINVAL     Invalid *queue\_selector* argument.
- ENOSYS    The resource manager associated with *fildev* doesn't support this call.
- ENOTTY    The argument *fildev* doesn't refer to a terminal device.

**Examples:**

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 int fildev;

 fildev = open("/dev/ser1", O_RDWR);

 /* Throw away all input data */
 tcflush(fildev, TCIFLUSH);

 close(fildev);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*tcdrain()*, *tcflow()*, *tcsendbreak()*

## ***tcgetattr()***

© 2004, QNX Software Systems Ltd.

*Get the current terminal control settings*

### **Synopsis:**

```
#include <termios.h>

int tcgetattr(int fdes,
 struct termios *termios_p);
```

### **Arguments:**

*fdes*            The file descriptor associated with the terminal device.

*termios\_p*      A pointer to a **termios** structure in which *tcgetattr()* can store the terminal's control attributes.

### **Library:**

**libc**

### **Description:**

The *tcgetattr()* function gets the current terminal control settings for the opened device indicated by *fdes*, and stores the results in the structure pointed to by *termios\_p*.

### **Returns:**

0      Success.

-1     An error occurred; *errno* is set.

### **Errors:**

EBADF      The *fdes* argument is invalid.

ENOSYS     The resource manager associated with *fdes* doesn't support this call.

ENOTTY     The *fdes* argument doesn't refer to a terminal device.

## Examples:

See *tcsetattr()*.

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *fpathconf()*, *tcsetattr()*, **termios**

Chapter 7 of POSIX 1003.1

## ***tcgetpgrp()***

© 2004, QNX Software Systems Ltd.

*Get the process group ID associated with a device*

### **Synopsis:**

```
#include <sys/types.h>
#include <unistd.h>

pid_t tcgetpgrp(int fildes);
```

### **Arguments:**

*fildes* A file descriptor that's associated with the device whose process group ID you want to get.

### **Library:**

libc

### **Description:**

The *tcgetpgrp()* function returns the process group ID of the foreground process that's associated with the device indicated by *fildes*.

### **Returns:**

The ID of foreground process group. If an error occurs, -1 is returned, and *errno* is set.

### **Errors:**

|        |                                                                               |
|--------|-------------------------------------------------------------------------------|
| EBADF  | The argument <i>fildes</i> is invalid.                                        |
| ENOSYS | The resource manager associated with <i>fildes</i> doesn't support this call. |
| ENOTTY | The argument <i>fildes</i> isn't associated with a terminal device.           |



## Examples:

```
#include <sys/types.h>
#include <unistd.h>
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 printf("STDIN directs breaks to pgrp %d\n",
 tcgetpgrp(0));
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*signal()*, *tcsetpgrp()*

## ***tcgetsid()***

© 2004, QNX Software Systems Ltd.

*Get the process group ID of the session leader for a controlling terminal*

### **Synopsis:**

```
#include <unistd.h>

int tcgetsid(int filedes);
```

### **Arguments:**

*filedes* A file descriptor that's associated with the device whose ID you want to get.

### **Library:**

libc

### **Description:**

The *tcgetsid()* function returns the process group ID of the session for which the terminal specified by *filedes* is the controlling terminal.

### **Returns:**

The process group ID associated with the terminal, or -1 if an error occurs (*errno* is set).

### **Errors:**

|        |                                                                           |
|--------|---------------------------------------------------------------------------|
| EACCES | The <i>filedes</i> argument isn't associated with a controlling terminal. |
| EBADF  | The <i>filedes</i> argument isn't a valid file descriptor.                |
| ENOTTY | The file associated with <i>filedes</i> isn't a terminal.                 |

### **Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*tcsetsid()*

## ***tcgetsize()***

© 2004, QNX Software Systems Ltd.

*Get the size of a character device*

---

### **Synopsis:**

```
#include <termios.h>

int tcgetsize(int filedes,
 int* prows,
 int* pcols);
```

### **Arguments:**

|                             |                                                                                             |
|-----------------------------|---------------------------------------------------------------------------------------------|
| <i>filedes</i>              | A file descriptor that's associated with the device whose size you want to get.             |
| <i>prows</i> , <i>pcols</i> | NULL, or pointers to locations where the function can store the number of rows and columns. |

### **Library:**

libc

### **Description:**

The *tcgetsize()* function gets the size of the character device associated with *filedes* and stores the number of rows and columns in *prows* and *pcols* if they're not NULL.

### **Returns:**

|    |                                           |
|----|-------------------------------------------|
| 0  | Success.                                  |
| -1 | An error occurred ( <i>errno</i> is set). |

### **Errors:**

|        |                                                                           |
|--------|---------------------------------------------------------------------------|
| EACCES | The <i>filedes</i> argument isn't associated with a controlling terminal. |
| EBADF  | The <i>filedes</i> argument isn't a valid file descriptor.                |
| ENOTTY | The file associated with <i>filedes</i> isn't a terminal.                 |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*tcsetsize()*

## ***tcinject()***

© 2004, QNX Software Systems Ltd.

*Inject characters into a device's input buffer*

### **Synopsis:**

```
#include <termios.h>

int tcinject(int fd,
 char *buf,
 int n);
```

### **Arguments:**

- fd* A file descriptor that's associated with the device whose input buffer you want to add characters to.
- buf* A pointer to a buffer that contains the characters that you want to insert.
- n* The number of characters to insert. If *n* is positive, the characters are written to the canonical (edited) queue. If *n* is negative, the characters are written to the raw queue.

### **Library:**

`libc`

### **Description:**

The *tcinject()* function injects *n* characters pointed to by *buf* into the input buffer of the device given in *fd*.

Note that while injecting into the canonical queue, editing characters in *buf* are acted upon as though the user entered them directly from the device. If *buf* doesn't contain a newline (`'\n'`), carriage return (`'\r'`) or a data-forwarding character such as an EOF, data doesn't become available for reading. If *buf* does contain a data-forwarding character, it should contain only one as the last character in *buf*.

This function is useful for implementing command-line recall algorithms by injecting recalled lines into the canonical queue.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EBADF The *fd* argument is invalid or the file isn't opened for read.
- ENOSYS This function isn't supported for the device opened.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <termios.h>

int main(void)
{
 char *p = "echo Hello world!\n";

 /* Inject the line all at once */
 tcinject(0, p, strlen(p));

 /* Inject the line one character at a time */
 while(*p)
 tcinject(0, p++, 1);
 return EXIT_SUCCESS;
}
```

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |

*continued...*

**Safety**

---

|        |     |
|--------|-----|
| Thread | Yes |
|--------|-----|



**Synopsis:**

```
#include <termios.h>

int tcischars(int filedes);
```

**Arguments:**

*filedes*     A file descriptor that's associated with the device that you want to check.

**Library:**

libc

**Description:**

The *tcischars()* function checks to see how many characters are waiting to be read from the given file descriptor, *filedes*.

**Returns:**

The number of characters waiting to be read, or -1 if an error occurred.

**Errors:**

ENOTTY     The *fd* argument isn't the file descriptor for a character device.

**Classification:**

QNX Neutrino

**Safety**

---

Cancellation point    No

Interrupt handler      No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Synopsis:**

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(AF_INET,
 SOCK_STREAM,
 0);
```

**Description:**

The TCP protocol provides reliable, flow-controlled, two-way transmission of data. It's a byte-stream protocol used to support the SOCK\_STREAM abstraction.

TCP uses the standard Internet address format and also provides a per-host collection of "port addresses." Thus, each address is composed of an Internet address specifying the host and network, with a specific TCP port on the host identifying the peer entity.

Sockets using the TCP protocol are either *active* or *passive*. Active sockets initiate connections to passive sockets. By default, TCP sockets are created active.

To create a passive socket, you must bind the socket with the *bind()* system call, and then use the *listen()* system call. Only passive sockets may use the *accept()* call to accept incoming connections; only active sockets may use the *connect()* call to initiate connections.

Passive sockets may "underspecify" their location to match incoming connection requests from multiple networks. With this technique, termed *wildcard addressing*, a single server can provide service to clients on multiple networks. If you wish to create a socket that listens on all networks, the Internet address INADDR\_ANY must be bound. You can still specify the TCP port at this time. If the port isn't specified, the system assigns one.

Once a connection has been established, the socket's address is fixed by the peer entity's location. The address assigned to the socket is the address associated with the network interface through which packets

are being transmitted and received. Normally this address corresponds to the peer entity's network.

TCP supports several socket options (defined in `<netinet/tcp.h>`) that you can set with `setsockopt()` and retrieve with `getsockopt()`. The option level for these calls is the protocol number for TCP, available from `getprotobyname()`.

#### TCP\_NODELAY

Under most circumstances, TCP sends data when it's presented. When outstanding data hasn't yet been acknowledged, TCP gathers small amounts of output to be sent in a single packet once an acknowledgment is received.

For a few clients (such as windowing systems that send a stream of mouse events that receive no replies), this packetization may cause significant delays. Therefore, TCP provides a boolean option, `TCP_NODELAY`, to defeat this algorithm.

#### TCP\_MAXSEG

The Maximum Segment Size (MSS) for a TCP connection. The value returned is the maximum amount of data that TCP sends to the other end. If this value is fetched before the socket is connected, the value returned is the default value that's used if an MSS option isn't received from the other end.

#### TCP\_KEEPAIVE

Specifies the idle time in seconds for the connection before TCP starts sending "keepalive" probes. The default value is 2 hours. This option is effective only when the `SO_KEEPAIVE` socket option is enabled.

You can use options at the IP transport level with TCP (see the IP protocol. Incoming connection requests that are source-routed are noted, and the reverse source route is used in responding.

## Returns:

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

## Errors:

|               |                                                                                                                   |
|---------------|-------------------------------------------------------------------------------------------------------------------|
| EADDRINUSE    | You tried to create a socket with a port that's already been allocated.                                           |
| EADDRNOTAVAIL | You tried to create a socket with a network address for which no network interface exists.                        |
| ECONNREFUSED  | The remote peer actively refused connection establishment (usually because no process was listening to the port). |
| ECONNRESET    | The remote peer forced the connection to be closed.                                                               |
| EISCONN       | You tried to establish a connection on a socket that already has one.                                             |
| ENOBUFS       | The system ran out of memory for an internal data structure.                                                      |
| ETIMEDOUT     | A connection was dropped due to excessive retransmissions.                                                        |

## See also:

IP protocol

*accept()*, *bind()*, *connect()*, *getprotobyname()*, *getsockopt()*, *listen()*, *setsockopt()*, *socket()*

*RFC 793*

## ***tcsendbreak()***

© 2004, QNX Software Systems Ltd.

*Assert a break condition over a communications line*

### **Synopsis:**

```
#include <termios.h>

int tcsendbreak(int fdes,
 int duration);
```

### **Arguments:**

*fd*            A file descriptor that's associated with the line that you want to assert the break on.

*duration*     The number of milliseconds that you want to break for.

### **Library:**

`libc`

### **Description:**

The *tcsendbreak()* function asserts a break condition over the communication line associated with the opened file descriptor indicated by *fdes*.

The break condition lasts for at least *duration* milliseconds, or approximately 300 milliseconds if *duration* is zero. The system rounds the effective value of *duration* up to the next highest supported interval, which is typically a multiple of 100 milliseconds.

### **Returns:**

0            Success.

-1          An error occurred (*errno* is set).

### **Errors:**

EBADF        The argument *fdes* is invalid.

ENOSYS      The resource manager associated with *fdes* doesn't support this call.

ENOTTY      The argument *fildev* doesn't refer to a terminal device.

## Examples:

```
#include <termios.h>
#include <fcntl.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 int fd;

 fd = open("/dev/ser1", O_RDWR);

 /* Send a 500 millisecond break */
 tcsendbreak(fd, 500);

 close(fd);
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*tcdrain()*, *tcflow()*, *tcflush()*

## ***tcsetattr()***

© 2004, QNX Software Systems Ltd.

*Change the terminal control settings for a device*

### **Synopsis:**

```
#include <termios.h>

int tcsetattr(int fildev,
 int optional_actions,
 const struct termios *termios_p);
```

### **Arguments:**

*fildev*            The file descriptor associated with the terminal device.

*termios\_p*        A pointer to a **termios** structure that describes the attributes that you want to set for the terminal device.

### **Library:**

**libc**

### **Description:**

The *tcsetattr()* function sets the current terminal control settings for the opened device indicated by *fildev* to the values stored in the structure pointed to by *termios\_p*.

The operation of *tcsetattr()* depends on the values in *optional\_actions*:

|           |                                                                                                                                         |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------|
| TCSANOW   | The change is made immediately.                                                                                                         |
| TCSADRAIN | No change is made until all currently written data has been transmitted.                                                                |
| TCSAFLUSH | No change is made until all currently written data has been transmitted, at which point any received but unread data is also discarded. |

The **termios** control structure is defined in **<termios.h>**. For more information, see *tcgetattr()*.



## Returns:

- 0 Success.
- 1 An error occurred; *errno* is set.

## Errors:

- EBADF The argument *fildev* is invalid;
- EINVAL The argument *action* is invalid, or one of the members of *termios\_p* is invalid.
- ENOSYS The resource manager associated with *fildev* doesn't support this call.
- ENOTTY The argument *fildev* doesn't refer to a terminal device.

## Examples:

```
#include <stdlib.h>
#include <termios.h>

int main(void)
{
 raw(0);
 /*
 * Stdin is now "raw"
 */
 unraw (0);
 return EXIT_SUCCESS;
}

int raw(fd)
int fd;
{
 struct termios termios_p;

 if(tcgetattr(fd, &termios_p))
 return(-1);

 termios_p.c_cc[VMIN] = 1;
 termios_p.c_cc[VTIME] = 0;
 termios_p.c_lflag &= ~(ECHO|ICANON|ISIG|
 ECHOE|ECHOK|ECHONL);
 termios_p.c_oflag &= ~(OPOST);
 return(tcsetattr(fd, TCSADRAIN, &termios_p));
}
```

```
 }

int unraw(fd)
int fd;
{
 struct termios termios_p;

 if(tcgetattr(fd, &termios_p))
 return(-1);

 termios_p.c_lflag |= (ECHO|ICANON|ISIG|
 ECHOE|ECHOK|ECHONL);
 termios_p.c_oflag |= (OPOST);
 return(tcsetattr(fd, TCSADRAIN, &termios_p));
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:***errno*, *select()*, *tcgetattr()*, **termios**

## Synopsis:

```
#include <sys/types.h>
#include <unistd.h>

int tcsetpgrp(int fildev,
 pid_t pgrp_id);
```

## Arguments:

*fildev*      A file descriptor that's associated with the device whose process group ID you want to set.

*pgrp\_id*     The process group ID that you want to assign to the device.

## Library:

libc

## Description:

The *tcsetpgrp()* function sets the process group ID associated with the device indicated by *fildev* to be *pgrp\_id*.

If successful, the *tcsetpgrp()* function causes subsequent breaks on the indicated terminal device to generate a SIGINT on all process in the given process group.

## Returns:

0      Success.

-1     An error occurred (*errno* is set).

## Errors:

EBADF      The argument *fildev* is invalid.

EINVAL     The argument *pgrp\_id* is invalid.

- ENOSYS      The resource manager associated with *fildev* doesn't support this call.
- ENOTTY      The argument *fildev* isn't associated with a terminal device.
- EPERM      The argument *pgrp\_id* isn't part of the same session as the calling process.

## Examples:

```
#include <sys/types.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 /*
 * Direct breaks on stdin to me
 */
 tcsetpgrp(0, getpid());
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*signal()*, *tcgetpgrp()*

## ***tcsetSID()***

© 2004, QNX Software Systems Ltd.

*Make a terminal device a controlling device*

### **Synopsis:**

```
#include <termios.h>

int tcsetSID(int fd,
 pid_t pid);
```

### **Arguments:**

*fd*      A file descriptor that's associated with the device that you want to make a controlling device.

*pid*     The ID of the process that you want to associate with the controlling device.

### **Library:**

`libc`

### **Description:**

The *tcsetSID()* function makes the terminal device associated with the file descriptor argument *fd* into a controlling terminal that's associated with the process *pid*. If successful, this call causes subsequent hangup conditions on the terminal device *fd* to generate a SIGHUP signal on the given process.

This call is equivalent to calling `ioctl( fd, TIOCSCTTY )` to set the controlling terminal to the current process. You can clear the controlling terminal by passing -1 as *fd*.

### **Returns:**

0        Success.

-1      Failure; *errno* is set.

## Errors:

- EBADF      Invalid file descriptor.
- EINVAL     The argument *pid* is invalid.
- ENOSYS, ENOTTY  
            The argument *fd* isn't associated with a terminal device.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*ioctl()*, *tcgetattr()*

## ***tcsetattr()***

© 2004, QNX Software Systems Ltd.

*Set the size of a character device*

### **Synopsis:**

```
#include <termios.h>

int tcsetattr(int filedes,
 int rows,
 int cols);
```

### **Arguments:**

*filedes*      A file descriptor that's associated with the device whose size you want to set.

*rows, cols*    The number of rows and columns that you want to use.

### **Library:**

libc

### **Description:**

The *tcsetattr()* function sets the size of the character device associated with *filedes* to the given number of rows and columns.

### **Returns:**

0      Success.

-1     An error occurred (*errno* is set).

### **Errors:**

EACCES      The *filedes* argument isn't associated with a controlling terminal.

EBADF      The *filedes* argument isn't a valid file descriptor.

EINVAL      The *rows* or *cols* argument is invalid.

ENOTTY      The file associated with *filedes* isn't a terminal.



**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*tcgetattr()*

## ***tell(), tell64()***

© 2004, QNX Software Systems Ltd.

*Determine the current file position*

---

### **Synopsis:**

```
#include <unistd.h>

off_t tell(int filedes);

off64_t tell(int filedes);
```

### **Arguments:**

*filedes*     The file descriptor of the file whose position you want to get.

### **Library:**

libc

### **Description:**

The *tell()* function determines the current file position for any subsequent *read()* or *write()* operation (that is, any subsequent unbuffered file operation). The *filedes* value is the file descriptor returned by a successful call to *open()*.

You can use the returned value in conjunction with *lseek()* to reset the current file position.

### **Returns:**

The current file position, expressed as the number of bytes from the start of the file, or -1 if an error occurs (*errno* is set). A value of 0 indicates the start of the file.

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <sys/stat.h>
#include <unistd.h>
#include <fcntl.h>

char buffer[]
```

```
 = { "A text record to be written" };

int main(void)
{
 int filedes ;
 int size_written;

 /* open a file for output */
 /* replace existing file if it exists */
 filedes = open("file",
 O_WRONLY | O_CREAT | O_TRUNC,
 S_IRUSR | S_IWUSR | S_IRGRP | S_IWGRP);

 if(filedes != -1) {

 /* print file position */
 printf("%ld\n", tell(filedes));

 /* write the text */
 size_written = write(filedes , buffer,
 sizeof(buffer));

 /* print file position */
 printf("%ld\n", tell(filedes));

 /* close the file */
 close(filedes);
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
0
28
```

## Classification:

*tell()* is QNX 4; *tell64()* is for large-file support

### Safety

Cancellation point Yes

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*chsize(), close(), creat(), dup(), dup2(), eof(), errno, execl(), execlp(),  
execlpe(), execv(), execve(), execvp(), execvpe(), fcntl(),  
fileno(), fstat(), isatty(), lseek(), open(), read(), sopen(), stat(),  
umask(), write()*

**Synopsis:**

```
#include <dirent.h>

long int telldir(DIR * dirp);
```

**Arguments:**

*dirp*      The directory stream for which you want to get the current location.

**Library:**

libc

**Description:**

The *telldir()* function obtains the current location associated with the directory stream specified by *dirp*.

**Returns:**

The current position of the specified directory stream, or -1 if an error occurs (*errno* is set).

**Errors:**

EBADF      The *dirp* argument doesn't refer to an open directory stream.

**Classification:**

POSIX 1003.1

**Safety**

---

Cancellation point    Yes

Interrupt handler      No

*continued...*

**Safety**

---

|                |    |
|----------------|----|
| Signal handler | No |
| Thread         | No |

**See also:**

*closedir()*, *errno*, *lstat()*, *opendir()*, *readdir()*, *readdir\_r()*, *rewinddir()*, *seekdir()*, *stat()*

## Synopsis:

```
#include <stdio.h>

char* tempnam(const char* dir,
 const char* pfx);
```

## Arguments:

*dir* NULL, or the directory to use in the pathname.

*pfx* NULL, or a prefix to use in the pathname.



---

If *pfx* isn't NULL, the string it points to must be no more than 5 bytes long.

---

## Library:

libc

## Description:

The *tempnam()* function generates a pathname for use as a temporary file. The pathname is in the directory specified by *dir* and has the prefix specified in *pfx*.

If *dir* is NULL, the pathname is prefixed with the first accessible directory contained in:

- the temporary file directory *P\_tmpdir* (defined in `<stdio.h>`)
- the **TMPDIR** environment variable
- the `_PATH_TMP` constant (defined in `<paths.h>`).

If all of these paths are inaccessible, *tempnam()* attempts to use `/tmp` and then the current working directory.

The *tempnam()* function generates up to `TMP_MAX` unique file names before it starts to recycle them.

**Returns:**

A pointer to the generated file name, which you should deallocate with the *free()* function when the application no longer needs it, or NULL if an error occurs.

**Errors:**

ENOMEM      There's insufficient memory available to create the pathname.

**Classification:**

Standard Unix

**Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

**Caveats:**

The *tempnam()* functions creates only pathnames; the application must create and remove the files.

It's possible for another thread or process to create a file with the same name between the time the pathname is created and the file is opened.

**See also:**

*free()*, *tmpfile()*, *tmpnam()*, *unlink()*



**Synopsis:**

```
struct termios {
 tcflag_t c_iflag;
 tcflag_t c_oflag;
 tcflag_t c_cflag;
 tcflag_t c_lflag;
 cc_t c_cc [NCCS];
 uint32_t reserved [3];
 speed_t c_ispeed;
 speed_t c_ospeed;
};
```

**Description:**

The **termios** control structure is defined in `<termios.h>`, and contains at least the members described below.

**tcflag\_t c\_iflag**

Input modes. This member contains at least the following bits:

|        |                                                             |
|--------|-------------------------------------------------------------|
| BRKINT | Signal interrupt on break.                                  |
| ICRNL  | Map CR to NL on input.                                      |
| IGNBRK | Ignore break conditions.                                    |
| IGNCR  | Ignore CR.                                                  |
| IGNPAR | Ignore characters with parity errors.                       |
| INLCR  | Map NL to CR on input.                                      |
| INPCK  | Enable input parity check.                                  |
| ISTRIP | Strip top bit from character.                               |
| IXOFF  | Enable software input flow control (via START/STOP chars).  |
| IXON   | Enable software output flow control (via START/STOP chars). |
| PARMRK | Mark parity errors in the input data stream.                |

**tcflag\_t** *c\_oflag*

Output modes. This member contains at least the following bits:

OPOST Perform output processing.

**tcflag\_t** *c\_cflag*

Control modes. This member contains at least the following bits:

CLOCAL Ignore modem status lines.

CREAD Enable receiver.

CSIZE Number of data bits per character.

CS5 5 data bits.

CS6 6 data bits.

CS7 7 data bits.

CS8 8 data bits.

CSTOPB Two stop bits, else one.

HUPCL Hang up on last close.

IHFLOW Support input flow control using the hardware handshaking lines.

OHFLOW Support output flow control using the hardware handshaking lines.

PARENB Parity enable.

PARODD Odd parity, else even.

PARSTK Stick parity (mark parity if PARODD is set, else space parity).

**tcflag\_t** *c\_lflag*

Local modes. This member contains at least the following bits:

ECHO Enable echo.

ECHOE Echo ERASE as destructive backspace.

ECHOK Echo KILL as a line erase.

ECHONL Echo '\n', even if ECHO is off.  
ICANON Canonical input mode (line editing enabled).  
IEXTEN QNX extensions to POSIX are enabled.  
ISIG Enable signals.  
NOFLSH Disable flush after interrupt, quit, or suspend.  
TOSTOP Send SIGTTOU for background output.

**cc\_t** *c\_cc*[NCCS]

Control characters. The array *c\_cc* includes at least the following control characters:

*c\_cc*[VEOF] EOF character.  
*c\_cc*[VEOL] EOL character.  
*c\_cc*[VERASE] ERASE character.  
*c\_cc*[VINTR] INTR character.  
*c\_cc*[VKILL] KILL character.  
*c\_cc*[VMIN] MIN value.  
*c\_cc*[VQUIT] QUIT character.  
*c\_cc*[VSUSP] SUSP character.  
*c\_cc*[VTIME] TIME value.  
*c\_cc*[VSTART] START character.  
*c\_cc*[VSTOP] STOP character.

The following control characters are also defined, but are only acted on if they're immediately preceded by the nonzero characters in *c\_cc*[VPREFIX][4], and are immediately followed by the nonzero characters in *c\_cc*[VSUFFIX][4] and the IEXTEN bit of *c\_lflag* is set:

*c\_cc*[VLEFT] Left cursor motion.  
*c\_cc*[VRIGHT] Right cursor motion.  
*c\_cc*[VUP] Up cursor motion.

`c_cc[VDOWN]` Down cursor motion.  
`c_cc[VINS]` Insert character.  
`c_cc[VDEL]` Delete character.  
`c_cc[VRUB]` Rubout character.  
`c_cc[VCAN]` Cancel character.  
`c_cc[VHOME]` Home character.  
`c_cc[VEND]` End character.

Any of the control characters in the `c_cc` array can be disabled by setting that character to the `_PC_VDISABLE` parameter which is returned by `fpathconf()` (typically a zero).

**`speed_t c_ispeed`**

Input baud rate. This member should be queried and set with the `cfgetispeed()` and `cfsetispeed()` functions.

**`speed_t c_ospeed`**

Output baud rate. This member should be queried and set with the `cfgetospeed()` and `cfsetospeed()` functions.

## Classification:

POSIX 1003.1

## See also:

`cfgetispeed()`, `cfgetospeed()`, `cfsetispeed()`, `cfsetospeed()`,  
`cfmakeraw()`, `fpathconf()`, `forkpty()`, `openpty()`, `pathconf()`,  
`readcond()`, `tcgetattr()`, `tcsetattr()`

**Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

thread_pool_control(thread_pool_t * pool,
 thread_pool_attr_t * attr,
 uint16_t lower,
 uint16_t upper,
 unsigned flags)
```

**Arguments:**

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>pool</i>         | A thread pool handle that was returned by <i>thread_pool_create()</i> .                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <i>attr</i>         | A pointer to a <b>thread_pool_attr_t</b> structure that specifies the attributes that you want to use for the thread pool. For more information, see “Thread-pool attributes,” in the documentation for <i>thread_pool_create()</i> .                                                                                                                                                                                                                                                                        |
| <i>lower, upper</i> | This function blocks until the number of threads created is between the range of <i>upper</i> and <i>lower</i> , unless you set <b>THREAD_POOL_CONTROL_NONBLOCK</b> in <i>flags</i> .                                                                                                                                                                                                                                                                                                                        |
| <i>flags</i>        | Which attributes you want to change for the thread pool; any combination of the following bits: <ul style="list-style-type: none"><li>● <b>THREAD_POOL_CONTROL_HIWATER</b> — adjust the high-water value of the number of threads allowed in the thread pool.</li><li>● <b>THREAD_POOL_CONTROL_INCREMENT</b> — adjust the increment value of the number of threads.</li><li>● <b>THREAD_POOL_CONTROL_LOWATER</b> — adjust the low-water value of the number of threads allowed in the thread pool.</li></ul> |

- `THREAD_POOL_CONTROL_MAXIMUM` — adjust the maximum value of the number of threads allowed in the thread pool.
- `THREAD_POOL_CONTROL_NONBLOCK` — don't block while creating threads.

**Library:**

`libc`

**Description:**

Use *thread\_pool\_control()* to specify a thread pool's behavior and adjust its attributes.



---

Having several threads call this function with the same thread pool handle isn't recommended.

---

**Returns:**

-1 if an error occurs (*errno* is set).

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*thread\_pool\_destroy()*, *thread\_pool\_create()*, *thread\_pool\_limits()*,  
*thread\_pool\_start()*

## ***thread\_pool\_create()***

© 2004, QNX Software Systems Ltd.

*Create a thread pool handle*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

thread_pool_t * thread_pool_create (
 thread_pool_attr_t * pool_attr,
 unsigned flags);
```

### **Arguments:**

- pool\_attr*     A pointer to a **thread\_pool\_attr\_t** structure that specifies the attributes that you want to use for the thread pool. For more information, see “Thread-pool attributes,” below.
- flags*         Flags (defined in **<sys/dispatch.h>**) that affect what happens to the thread that’s creating the pool:
- **POOL\_FLAG\_EXIT\_SELF** — when the pool is started using *thread\_pool\_start()*, exit the thread that called this function.
  - **POOL\_FLAG\_USE\_SELF** — when the pool is started, use the calling thread as part of the pool.

### **Library:**

**libc**

### **Description:**

The *thread\_pool\_create()* function creates a thread pool handle. This handle is then used to start a thread pool with *thread\_pool\_start()*. With the thread pool functions, you can create and manage a pool of worker threads.



## How it works

The worker threads work in the following way:

- When a new worker thread is created, a context is allocated, which the thread uses to do its work.
- The thread then calls the blocking function. This function blocks until the thread has work to do. For example, the blocking function could call *MsgReceive()* to wait for a message.
- After the blocking function returns, the worker thread calls the handler function, which performs the actual work.
- When the handler function returns, the thread calls the blocking function again.

The thread continues to block and handle events until the thread pool decides this worker thread is no longer needed. Finally, when the worker thread exits, it releases the allocated context.

The thread pool manages these worker threads so that there's a certain number of them in the blocked state. Thus, as threads become busy in the handler function, the thread pool creates new threads to keep a minimum number of threads in a state where they can accept requests from clients. By the same token, if the demand on the thread pool goes down, the thread pool lets some of these blocked threads exit.

## Thread-pool attributes

The *pool\_attr* argument sets the:

- functions that get called, when a new thread is started or one dies, to allocate and free contexts used by threads
- blocking and handler functions
- parameters of the thread pool such as the number of worker threads, etc.

The *thread\_pool\_attr\_t* structure that the *pool\_attr* argument points to is defined as:

```
typedef struct _thread_pool_attr {
 THREAD_POOL_HANDLE_T *handle;
 THREAD_POOL_PARAM_T *(*block_func)
 (THREAD_POOL_PARAM_T *ctp);
 THREAD_POOL_PARAM_T *(*context_alloc)
 (THREAD_POOL_HANDLE_T *handle);
 void (*unblock_func)
 (THREAD_POOL_PARAM_T *ctp);
 int (*handler_func)
 (THREAD_POOL_PARAM_T *ctp);
 void (*context_free)
 (THREAD_POOL_PARAM_T *ctp);

 pthread_attr_t *attr;
 unsigned short lo_water;
 unsigned short increment;
 unsigned short hi_water;
 unsigned short maximum;
 unsigned reserved[8];
} thread_pool_attr_t;
```

The members include:

- |                      |                                                                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>handle</i>        | A handle that gets passed to the <i>context_alloc</i> function.                                                                                                                                                                                 |
| <i>block_func</i>    | The function that's called when the worker thread is ready to block, waiting for work. It returns a pointer that's passed to <i>handler_func</i> . If the function returns NULL, it indicates to the thread pool that this thread should exist. |
| <i>context_alloc</i> | The function that's called when a new thread is created by the thread pool. It is passed <i>handle</i> . The function returns a pointer, which is then passed to the blocking function <i>block_func</i> .                                      |
| <i>unblock_func</i>  | The function that's called to unblock threads. If you use <i>dispatch_block()</i> as the <i>block_func</i> , use <i>dispatch_unblock()</i> as the <i>unblock_func</i> .                                                                         |
| <i>handler_func</i>  | The function that's called after <i>block_func</i> returns to do some work. The function is passed in the pointer returned by <i>block_func</i> .                                                                                               |

|                     |                                                                                                                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>context_free</i> | The function that's called when the worker thread exits, to free the context allocated with <i>context_alloc</i> .                                                                                                                                                                                  |
| <i>attr</i>         | A pointer to a <i>pthread_attr_*</i> ( <i>pthread_attr_*</i> ) function that's passed to <i>pthread_create</i> ( <i>pthread_attr_*</i> ). The <i>pthread_attr_*</i> ( <i>pthread_attr_*</i> ) functions set the stack size, priority, etc. of the worker threads. If NULL, default values are used. |
| <i>lo_water</i>     | The minimum number of threads that the pool should keep in the blocked state (i.e. threads that are ready to do work).                                                                                                                                                                              |
| <i>increment</i>    | The number of new threads created at one time.                                                                                                                                                                                                                                                      |
| <i>hi_water</i>     | The maximum number of threads to keep in a blocked state.                                                                                                                                                                                                                                           |
| <i>maximum</i>      | The maximum number of threads that the pool can create.                                                                                                                                                                                                                                             |

## Returns:

A thread pool handle, or NULL if an error occurs (*errno* is set).

## Errors:

ENOMEM    Insufficient memory to allocate internal data structures.

## Examples:

Here's a simple multithreaded resource manager:

```
/* Define an appropriate interrupt number: */
#define INTNUM 0

#include <stdio.h>
#include <stddef.h>
#include <stdlib.h>
#include <sys/iofunc.h>
#include <sys/dispatch.h>
#include <sys/neutrino.h>
```

```
static resmgr_connect_funcs_t connect_funcs;
static resmgr_io_funcs_t io_funcs;
static iofunc_attr_t attr;

void *interrupt_thread(void *data)
/* *data isn't used */
{
 struct sigevent event;
 int id;

 /* fill in "event" structure */
 memset(&event, 0, sizeof(event));
 event.sigev_notify = SIGEV_INTR;

 /* INTNUM is the desired interrupt level */
 id = InterruptAttachEvent(INTNUM, &event, 0);

 :

 while (1) {
 InterruptWait(0, NULL);
 /*
 * do something about the interrupt,
 * perhaps updating some shared
 * structures in the resource manager
 */
 /* unmask the interrupt when done */
 InterruptUnmask(INTNUM, id);
 }
}

int
main(int argc, char **argv) {
 thread_pool_attr_t pool_attr;
 thread_pool_t *tpp;
 dispatch_t *dpp;
 resmgr_attr_t resmgr_attr;
 int id;

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr,
 "%s: Unable to allocate dispatch handle.\n",
 argv[0]);
 return EXIT_FAILURE;
 }
}
```

```

memset(&pool_attr, 0, sizeof pool_attr);
pool_attr.handle = dpp;
/* We are only doing resmgr-type attach */
pool_attr.context_alloc = resmgr_context_alloc;
pool_attr.block_func = resmgr_block;
pool_attr.handler_func = resmgr_handler;
pool_attr.context_free = resmgr_context_free;
pool_attr.lo_water = 2;
pool_attr.hi_water = 4;
pool_attr.increment = 1;
pool_attr.maximum = 50;

if((tpp = thread_pool_create(&pool_attr,
 POOL_FLAG_EXIT_SELF)) == NULL) {
 fprintf(stderr,
 "%s: Unable to initialize thread pool.\n",
 argv[0]);
 return EXIT_FAILURE;
}

iofunc_func_init(_RESMGR_CONNECT_NFUNCS,
 &connect_funcs,
 _RESMGR_IO_NFUNCS, &io_funcs);
iofunc_attr_init(&attr, S_IFNAM | 0666, 0, 0);

memset(&resmgr_attr, 0, sizeof resmgr_attr);
resmgr_attr.nparts_max = 1;
resmgr_attr.msg_max_size = 2048;

if((id = resmgr_attach(dpp, &resmgr_attr,
 "/dev/mynull",
 _FTYPE_ANY, 0, &connect_funcs,
 &io_funcs,
 &attr)) == -1) {
 fprintf(stderr,
 "%s: Unable to attach name.\n", argv[0]);
 return EXIT_FAILURE;
}

/* Start the thread which will handle interrupt events. */
pthread_create (NULL, NULL, interrupt_thread, NULL);

/* Never returns */
thread_pool_start(tpp);
}

```

For more examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, and *resmgr\_attach()*. For information on advanced topics in designing and implementing a resource manager, see

“Combine messages” section of the Writing a Resource Manager chapter in the *Programmer’s Guide*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*dispatch\_block()*, *dispatch\_create()*, *dispatch\_unblock()*,  
*pthread\_create()*, *resmgr\_attach()*, *select\_attach()*,  
*thread\_pool\_destroy()*, *thread\_pool\_start()*

## ***thread\_pool\_destroy()***

*Free the memory allocated to a thread pool*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int thread_pool_destroy(thread_pool_t * pool);
```

### **Arguments:**

*pool*     A thread pool handle that was returned by *thread\_pool\_create()*.

### **Library:**

libc

### **Description:**

The *thread\_pool\_destroy()* function frees the memory allocated to a thread pool that's identified by the handle *pool*. This is done only after all the threads of the thread pool have exited.



Prior to QNX Neutrino 6.1.0, this function simply deallocated the thread pool handle and returned. Although this was acceptable for servers that never exited (and consequently never shut down their thread pools), it's unsuitable for closing down a thread pool.

---

The *thread\_pool\_destroy()* function calls the unblock handler provided in the pool attribute structure. The unblock handler is called at least once for every thread in the thread pool. Once the unblock handler is called, the thread calling *thread\_pool\_destroy()* blocks until the number of threads in the thread pool drops to zero. When there are no more threads in the thread pool, the handle pool is freed and *thread\_pool\_destroy()* returns.



---

A side effect of this behavior is that a thread that's created by the thread pool can't call *thread\_pool\_destroy()* because the thread pool count will never drop to zero, and subsequently the function will never return.

---

## Returns:

- 0 Success.
- 1 An error occurred.

## Examples:

```
#include <sys/dispatch.h>
#include <stdio.h>

int main(int argc, char **argv) {
 thread_pool_t *tpp;

 :

 thread_pool_destroy (tpp);
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |



**See also:**

*thread\_pool\_control()*, *thread\_pool\_create()*, *thread\_pool\_limits()*,  
*thread\_pool\_start()*

## ***thread\_pool\_limits()***

© 2004, QNX Software Systems Ltd.

*Convenience wrapper function for `thread_pool_control()`*

### **Synopsis:**

```
#include <sys/iofunc.h>
#include <sys/dispatch.h>

int thread_pool_limits(thread_pool_t * pool,
 int lowater,
 int hiwater,
 int maximum,
 int increment,
 unsigned flags);
```

### **Arguments:**

|                  |                                                                                                                                                                                            |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>pool</i>      | A thread pool handle that was returned by <i>thread_pool_create()</i> .                                                                                                                    |
| <i>lowater</i>   | The minimum number of threads that the pool should keep in the blocked state (i.e. threads that are ready to do work), or a negative number if you don't want to change the current value. |
| <i>hiwater</i>   | The maximum number of threads that the pool should keep in the blocked state, or a negative number if you don't want to change the current value.                                          |
| <i>maximum</i>   | The maximum number of threads that the pool can create, or a negative number if you don't want to change the current value.                                                                |
| <i>increment</i> | The number of new threads created at one time, or a negative number if you don't want to change the current value.                                                                         |
| <i>flags</i>     | The only flag that's accepted is <code>THREAD_POOL_CONTROL_NONBLOCK</code> . For more information, see the documentation for <i>thread_pool_control()</i> .                                |

**Library:**

libc

**Description:**

The *thread\_pool\_limits()* function is a wrapper function for *thread\_pool\_control()*. If the value of *lowater*, *hiwater*, *maximum* or *increment* is  $\geq 0$  then that value is adjusted in the thread pool according to the handle *pool*.

If you don't set `THREAD_POOL_CONTROL_NONBLOCK`, the upper and lower bounds for waiting are:

- `lower = (lowater != -1) : lowater ? 0;`
- `upper = (maximum != -1) : maximum ? USHRT_MAX;`



Having several threads call this function with the same thread pool handle isn't recommended.

---

**Returns:**

-1 if an error occurs (*errno* is set).

**Classification:**

QNX Neutrino

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## *thread\_pool\_limits()*

© 2004, QNX Software Systems Ltd.

---

### **See also:**

*thread\_pool\_control()*, *thread\_pool\_create()*, *thread\_pool\_destroy()*,  
*thread\_pool\_start()*

**Synopsis:**

```
#include <sys/dispatch.h>

int thread_pool_start(void *pool);
```

**Arguments:**

*pool*     A thread pool handle that was returned by *thread\_pool\_create()*.

**Library:**

libc

**Description:**

The *thread\_pool\_start()* function starts the thread pool *pool*. The function may or may not return, depending on the flags that you passed to *thread\_pool\_create()*.

**Returns:**

EOK     Success.

-1     An error occurred.

**Examples:**

```
#include <sys/dispatch.h>
#include <stdio.h>
#include <stdlib.h>

int main(int argc, char **argv) {
 thread_pool_attr_t pool_attr;
 thread_pool_t *tpp;
 dispatch_t *dpp;
 resmgr_attr_t attr;
 resmgr_context_t *ctp;

 if((dpp = dispatch_create()) == NULL) {
 fprintf(stderr, "%s: Unable to allocate \
 dispatch context.\n", argv[0]);
 }
}
```

```
 return EXIT_FAILURE;
 }

 memset(&pool_attr, 0, sizeof (pool_attr));
 pool_attr.handle = dpp;
 /* We are only doing resmgr-type attach */
 pool_attr.context_alloc = resmgr_context_alloc;
 pool_attr.block_func = resmgr_block;
 pool_attr.handler_func = resmgr_handler;
 pool_attr.context_free = resmgr_context_free;
 pool_attr.lo_water = 2;
 pool_attr.hi_water = 4;
 pool_attr.increment = 1;
 pool_attr.maximum = 50;

 if((tpp = thread_pool_create(&pool_attr,
 POOL_FLAG_EXIT_SELF)) == NULL) {
 fprintf(stderr, "%s: Unable to initialize \
 thread pool.\n", argv[0]);
 return EXIT_FAILURE;
 }

 :

 /* Never returns */
 thread_pool_start(tpp);
}
```

For examples using the dispatch interface, see *dispatch\_create()*, *message\_attach()*, *resmgr\_attach()*, and *thread\_pool\_create()*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*thread\_pool\_create(), thread\_pool\_destroy()*

## ***ThreadCancel()*, *ThreadCancel\_r()*** © 2004, QNX Software Systems Ltd.

*Cancel a thread*

### **Synopsis:**

```
#include <sys/neutrino.h>

int ThreadCancel(int tid,
 void (*canstub)(void));

int ThreadCancel_r(int tid,
 void (*canstub)(void));
```

### **Arguments:**

|                |                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------|
| <i>tid</i>     | The ID of the thread that you want to destroy, as returned by <i>ThreadCreate()</i> .                  |
| <i>canstub</i> | A pointer to the location that you want the thread to jump to when the cancellation occurs; see below. |



---

You must provide a *canstub* function.

---

### **Library:**

libc

### **Description:**

These kernel calls request that the thread specified by *tid* be canceled. The target thread's cancelability state and type determine when the cancellation takes effect.

The *ThreadCancel()* and *ThreadCancel\_r()* functions are identical, except in the way they indicate errors. See the Returns section for details.

When the cancellation is acted upon, the thread jumps to the location specified by *canstub*. This stub should call cancellation cleanup handlers for the thread. When the last cancellation cleanup handler returns, the stub must terminate the thread using:

```
ThreadDestroy(0, -1, PTHREAD_CANCEL);
```



Unlike *ThreadDestroy()*, which destroys a thread immediately, *ThreadCancel()* requests that the target thread execute any cleanup code and then terminate at its earliest convenience.

The cancellation processing in the target thread runs asynchronously with respect to the calling thread, which doesn't block.

The combinations of cancelability state and type are as follows:

| <b>State</b> | <b>Type</b> | <b>Description</b>                                                                                                                                                                                                                                                                                                                          |
|--------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disabled     | Deferred    | Cancel requests are made pending                                                                                                                                                                                                                                                                                                            |
| Disabled     | Async       | Cancel requests are made pending                                                                                                                                                                                                                                                                                                            |
| Enabled      | Deferred    | Cancellation happens at the next cancellation point. These are at explicitly coded calls to <i>pthread_testcancel()</i> or an attempt to enter a blocking state in any of the calls defined in the table below. All kernel calls that block are cancellation points, with the exception of <i>MsgSendvnc()</i> and <i>SyncMutexLock()</i> . |
| Enabled      | Async       | Cancellation happens immediately.                                                                                                                                                                                                                                                                                                           |

Use *pthread\_setcancelstate()*, and *pthread\_setcanceltype()* to set the state and type.

POSIX defines a list of functions that are cancellation points; some functions that aren't listed there may also be cancellation points. For a full list, see "Cancellation points" in the appendix, Summary of Safety Information. Any function that calls a blocking kernel call that's a cancellation point will itself become a cancellation point when the kernel call is made. The most common blocking kernel call in library code is *MsgSendv()*.

### Blocking states

These calls don't block.

### Returns:

The only difference between these functions is the way they indicate errors:

#### *ThreadCancel()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

#### *ThreadCancel\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

### Errors:

ESRCH     The thread indicated by *tid* doesn't exist.

### Classification:

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_setcancelstate(), pthread\_setcanceltype(),  
pthread\_testcancel(), ThreadCreate(), ThreadDestroy()*

# ***ThreadCreate()*, *ThreadCreate\_r()*** © 2004, QNX Software Systems Ltd.

*Create a thread*

## **Synopsis:**

```
#include <sys/neutrino.h>

int ThreadCreate(
 pid_t pid,
 void* (func) (void*),
 void* arg,
 const struct _thread_attr* attr);

int ThreadCreate_r(
 pid_t pid,
 void* (func) (void*),
 void* arg,
 const struct _thread_attr* attr);
```

## **Arguments:**

- pid* The ID of the process that you want to create the thread in, or 0 to create the thread in the current process.
- func* A pointer to the function that you want the thread to execute. The *arg* argument that you pass to *ThreadCreate()* is passed to *func()* as its sole argument. If *func()* returns, it returns to the address defined in the *exitfunc* member of *attr*.
- arg* A pointer to any data that you want to pass to *func*.
- attr* A pointer to a **\_thread\_attr** structure that specifies the attributes for the new thread, or NULL if you want to use the default attributes.



If you modify the attributes after creating the thread, the thread isn't affected.

---

For more information, see “Thread attributes,” below.

## Library:

`libc`

## Description:

These kernel calls create a new thread of execution, with attributes specified by *attr*, within the process specified by *pid*. If *pid* is zero, the current process is used.



---

Only the Process Manager can create threads in another process.

---

The *ThreadCreate()* and *ThreadCreate\_r()* functions are identical, except in the way they indicate errors. See the Returns section for details.

The new thread shares all resources of the process in which it's created. This includes memory, timers, channels and connections. The standard C library contains mutexes to make it thread-safe.

## Thread attributes

The `_thread_attr` structure pointed to by *attr* contains at least the following members:

`int flags` See below for a list of *flags*. The default flag is always zero.

`size_t stacksize`

The stack size of the thread stack defined in the *stackaddr* member. If *stackaddr* is NULL, then *stacksize* specifies the size of stack to dynamically allocate. If *stacksize* is zero, then 4096 bytes are assumed. The minimum allowed *stacksize* is defined by PTHREAD\_STACK\_MIN.

`void* stackaddr`

NULL, or the address of a stack that you want the thread to use. Set the *stacksize* member to the size of the stack.

If you provide a non-NULL *stackaddr*, it's your responsibility to release the stack when the thread dies. If *stackaddr* is NULL, then the kernel dynamically allocates a stack on thread creation and automatically releases it on the thread's death.

**void\*** (*exitfunc*)(**void\*** *status*)

The address to return to if the thread function returns.



---

The thread *returns* to *exitfunc*. This means that the *status* variable isn't passed as a normal parameter. Instead, it appears in the return-value position dictated by the CPU's calling convention (e.g. **EAX** on an x86, **R3** on PPC, **v0** on MIPS, and so on).

The *exitfunc* function normally has to have compiler- and CPU-specific manipulation to access the *status* data (pulling it from the return register location to a proper local variable). Alternatively, you can write the *exitfunc* function in assembly language for each CPU.

---

**int** *policy*     The scheduling policy, as defined by the *SchedSet()* kernel call. This member is used only if you set the **PTHREAD\_EXPLICIT\_SCHED** flag. If you want the thread to inherit the policy, but you want to specify the scheduling parameters in the *param* member, set the **PTHREAD\_EXPLICIT\_SCHED** flag and set the *policy* member to **SCHED\_NOCHANGE**.

**struct sched\_param** *param*

A **sched\_param** structure that specifies the scheduling parameters, as defined by the *SchedSet()* kernel call. This member is used only if you set the **PTHREAD\_EXPLICIT\_SCHED** flag.

You can set the *attr* argument's *flags* member to a combination of the following:

PTHREAD\_CREATE\_JOINABLE (default)

Put the thread into a zombie state when it terminates. It stays in this state until you retrieve its exit status or detach the thread.

PTHREAD\_CREATE\_DETACHED

Create the thread in the detached state; it doesn't become a zombie. You can't call *ThreadJoin()* for a detached thread.

PTHREAD\_INHERIT\_SCHED (default)

Use the scheduling attributes of the creating thread for the new thread.

PTHREAD\_EXPLICIT\_SCHED

Take the scheduling policy and parameters for the new thread from the *policy* and *param* members of *attr*.

PTHREAD\_SCOPE\_SYSTEM (default)

Schedule the thread is against all threads in the system.

PTHREAD\_SCOPE\_PROCESS

Don't set this flag; the QNX Neutrino OS implements true microkernel threads that have only a system scope.

PTHREAD\_MULTISIG\_ALLOW (default)

If the thread dies because of an unblocked, uncaught signal, terminate all threads, and hence, the process.

PTHREAD\_MULTISIG\_DISALLOW

Terminate only this thread; all other threads in the process are unaffected.

PTHREAD\_CANCEL\_DEFERRED (default)

Cancellation occurs only at cancellation points as defined by *ThreadCancel()*.

PTHREAD\_CANCEL\_ASYNCHRONOUS

Every opcode executed by the thread is considered a cancellation point. The POSIX and C library aren't asynchronous-cancel safe.

## Signal state

The signal state of the new thread is initialized as follows:

- The signal mask is inherited from the creating thread.
- The set of pending signals is empty.
- The cancel state and type are `PTHREAD_CANCEL_ENABLE` and `PTHREAD_CANCEL_DEFERRED`.

## Local storage for private data

Each thread contains a thread local storage area for its private data. This area can be accessed using the global variable `_TLS` defined in `<sys/neutrino.h>` as a pointer. The kernel ensures that `_TLS` always points to the thread local storage for the thread that's running.

The thread local storage is defined by the structure `_thread_local_storage`, which contains at least the following members:

|                                       |                                                                                                                                                                                                           |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>void* (exitfunc)(void *)</code> | The exit function to call if the thread returns.                                                                                                                                                          |
| <code>void* arg</code>                | The sole argument that was passed to the thread.                                                                                                                                                          |
| <code>int* errptr</code>              | A pointer to a thread unique <i>errno</i> value. For the main thread, this points to the global variable <i>errno</i> . For all other threads, this points to the member <i>errval</i> in this structure. |
| <code>int errval</code>               | A thread-unique <i>errno</i> that the thread uses if it isn't the main thread.                                                                                                                            |
| <code>int flags</code>                | The thread flags used on thread creation in addition to runtime flags used for implementing thread cancellation.                                                                                          |
| <code>pid_t pid</code>                | The ID of the process that contains the thread.                                                                                                                                                           |
| <code>int tid</code>                  | The thread's ID.                                                                                                                                                                                          |



## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *ThreadCreate()*

The thread ID of the newly created thread. If an error occurs, the function returns -1 and sets *errno*.

### *ThreadCreate\_r()*

The thread ID of the newly created thread. This function does **NOT** set *errno*. If an error occurs, the function returns the negative of a value from the Errors section.

## Errors:

|         |                                                                                                                                                                          |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN  | All kernel thread objects are in use.                                                                                                                                    |
| EFAULT  | A fault occurred when the kernel tried to access the buffers provided.                                                                                                   |
| EINVAL  | Invalid scheduling policy or priority specified.                                                                                                                         |
| ENOTSUP | PTHREAD_SCOPE_PROCESS was requested. All kernel threads are PTHREAD_SCOPE_SYSTEM.                                                                                        |
| EPERM   | The calling thread doesn't have sufficient permission to create a thread in another process. Only a thread with a process ID of 1 can create threads in other processes. |
| ESRCH   | The process indicated by <i>pid</i> doesn't exist.                                                                                                                       |

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## **Caveats:**

The QNX interpretation of PTHREAD\_STACK\_MIN is enough memory to run a thread that does nothing:

```
void nothingthread(void)
{
 return;
}
```

## **See also:**

*sched\_param, SchedSet(), ThreadCancel(), ThreadDestroy()*

**Synopsis:**

```
#include <sys/neutrino.h>

int ThreadCtl(int cmd,
 void * data);

int ThreadCtl_r(int cmd,
 void * data);
```

**Arguments:**

*cmd*      The command you want to execute; see below.

*data*     A pointer to data associated with the specific command; see below.

**Library:**

`libc`

**Description:**

These kernel calls allow you to make QNX-specific changes to a thread.

The *ThreadCtl()* and *ThreadCtl\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

The following calls are defined:

*ThreadCtl*(\_NTO\_TCTL\_ALIGN\_FAULT, *data*)

Control the misaligned access response. The *data* argument is a pointer to an `int` whose value indicates how you want to respond:

- Greater than 0 — make a misaligned access fault with a SIGBUS, if the architecture permits it.
- Less than 0 — make the kernel attempt to emulate an instruction with a misaligned access. If the attempt fails, it also faults with a SIGBUS.

- 0 — don't change the alignment-fault handling for the thread.

The function sets *data* to a positive or negative number, indicating the previous state of the the alignment-fault handling.

## *ThreadCtl*(\_NTO\_TCTL\_IO, 0)

Request I/O privity; let the thread execute the I/O opcodes *in*, *ins*, *out*, *outs*, *cli*, *sti* on architectures where it has the appropriate privilege, and let it attach IRQ handlers. You need *root* permissions to use this command. If a thread attempts to use these opcodes without successfully executing this call, the thread faults with a SIGSEGV when the opcode is attempted.



---

Threads created by the calling thread inherit the *\_NTO\_TCTL\_IO* status.

---

## *ThreadCtl*(\_NTO\_TCTL\_RUNMASK, (*int*)*runmask*)

Set processor affinity for the calling thread in a multiprocessor system. Each set bit in *runmask* represents a processor that the thread can run on.

By default, a thread's *runmask* is set to all ones, which allows it to run on any available processor. A value of *0x01* would, for example, force the thread to only run on the first processor.

You can use *\_NTO\_TCTL\_RUNMASK* to optimize the runtime performance of your system by, for example, relegating nonrealtime threads to a specific processor. In general, this shouldn't be necessary, since the QNX realtime scheduler always preempts a lower-priority thread immediately when a higher priority thread becomes ready.

The main effect of processor locking is the effectiveness of the CPU cache, since threads can be prevented from migrating.




---

Threads created by the calling thread don't inherit the `_NTO_TCTL_RUNMASK` status.

---

*ThreadCtl*(`_NTO_TCTL_THREADS_CONT`, 0)

Unfreeze all threads in the current process that were frozen using the `_NTO_TCTL_THREADS_HOLD` command.

*ThreadCtl*(`_NTO_TCTL_THREADS_HOLD`, 0)

Freeze all threads in the current process except the calling thread.

---




---

Threads created by the calling thread aren't frozen.

---

The *data* pointer is reserved for passing extra data for new commands that may be introduced in the future.

### Blocking states

These calls don't block.

### Returns:

The only difference between these functions is the way they indicate errors:

*ThreadCtl*()      If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

*ThreadCtl\_r*()      EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

### Errors:

EPERM      The process doesn't have superuser capabilities.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*InterruptDisable(), InterruptEnable(), InterruptMask(),  
InterruptUnmask()*

## ***ThreadDestroy(), ThreadDestroy\_r()***

*Destroy a thread immediately*

### **Synopsis:**

```
#include <sys/neutrino.h>

int ThreadDestroy(int tid,
 int priority,
 void* status);

int ThreadDestroy_r(int tid,
 int priority,
 void* status);
```

### **Arguments:**

*tid*            The ID of the thread that you want to destroy, as returned by *ThreadCreate()*, or 0 to destroy the current thread, or -1 to destroy all the threads in the current process.

*priority*       The priority at which you want to destroy multiple threads, or -1 to use the priority of the current thread.

*status*         The value to make available to a thread that joins a nondetached thread that's destroyed.

### **Library:**

`libc`

### **Description:**

These kernel calls terminate the thread specified by *tid*. If *tid* is 0, the calling thread is assumed. If *tid* is -1, all of the threads in the process are destroyed. When multiple threads are destroyed, the destruction is scheduled one thread at a time at the priority specified by the *priority* argument. If *priority* is -1, then the priority of the calling thread is used.

The *ThreadDestroy()* and *ThreadDestroy\_r()* functions are identical, except in the way they indicate errors. See the Returns section for details.

If a terminated thread isn't detached, it makes the value specified by the *status* argument available to any successful join on it. Until another thread retrieves this value, the thread ID *tid* isn't reused, and a small kernel resource (a thread object) is held in the system. If the thread is detached, then *status* is ignored, and all thread resources are immediately released.

When the last thread in a process is destroyed, the process terminates, and all thread resources are released, even if they're not detached and unjoined.



---

On return from *ThreadDestroy()* or *ThreadDestroy\_r()*, the target thread is marked for death, but if it isn't possible to kill it immediately, it may not be terminated until it attempts to run.

---

## Blocking states

If these calls return, they don't block.

## Returns:

If the calling thread is destroyed, *ThreadDestroy()* and *ThreadDestroy\_r()* don't return.

The only difference between these functions is the way they indicate errors:

### *ThreadDestroy()*

If this function returns and an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

### *ThreadDestroy\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If this function returns and an error occurs, any value in the Errors section may be returned.



**Errors:**

ESRCH    The thread indicated by *tid* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ThreadCancel(), ThreadCreate()*

# ***ThreadDetach()*, *ThreadDetach\_r()*** © 2004, QNX Software Systems Ltd.

*Detach a thread from a process*

## **Synopsis:**

```
#include <sys/neutrino.h>

int ThreadDetach(int tid);

int ThreadDetach_r(int tid);
```

## **Arguments:**

*tid* The ID of the thread that you want to detach, as returned by *ThreadCreate()*, or 0 to detach the current thread.

## **Library:**

libc

## **Description:**

These kernel calls detach the thread specified by *tid*. If *tid* is zero, the calling thread is used. Once detached, attempts to call *ThreadJoin()* on *tid* fail. When a detached thread terminates, its termination status is discarded and all its resources are released.

The *ThreadDetach()* and *ThreadDetach\_r()* functions are identical, except in the way they indicate errors. See the Returns section for details.

## **Blocking states**

These calls don't block.

## **Returns:**

The only difference between these functions is the way they indicate errors:

*ThreadDetach()*

If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

*ThreadDetach\_r()*

Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function can return any value listed in the Errors section.

**Errors:**

- EINVAL     The thread is already detached.
- ESRCH     The thread indicated by *tid* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ThreadCreate(), ThreadJoin()*

## ***ThreadJoin()*, *ThreadJoin\_r()***

© 2004, QNX Software Systems Ltd.

*Block until a thread terminates*

### **Synopsis:**

```
#include <sys/neutrino.h>

int ThreadJoin(int tid,
 void** status);

int ThreadJoin_r(int tid,
 void** status);
```

### **Arguments:**

*tid*            The ID of the thread that you want to detach, as returned by *ThreadCreate()*.

*status*        The address of a pointer to a location where the function can store the thread's exit status.

### **Library:**

`libc`

### **Description:**

The *ThreadJoin()* and *ThreadJoin\_r()* kernel calls block until the thread specified by *tid* terminates. If *status* isn't NULL, the functions save the thread's exit status in the area pointed to by *status*. If the thread *tid* has already terminated, the functions immediately return with success and the status, if requested.

These functions are identical except in the way they indicate errors. See the Returns section for details.

When *ThreadJoin()* returns successfully, the target thread has been successfully terminated. Until this occurs, the thread ID *tid* isn't reused and a small kernel resource (a thread object) is retained.

You can't join a thread that's detached (see *ThreadCreate()* and *ThreadDetach()*).

The target thread must be joinable. Multiple *pthread\_join()*, *pthread\_timedjoin()*, *ThreadJoin()*, and *ThreadJoin\_r()* calls on the same target thread aren't allowed.

**Blocking states**

STATE\_JOIN The calling thread blocks waiting for the indicated thread to exit.

**Returns:**

The only difference between these functions is the way they indicate errors:

*ThreadJoin()* If an error occurs, the function returns -1 and sets *errno*. Any other value returned indicates success.

*ThreadJoin\_r()* Returns EOK on success. This function does **NOT** set *errno*. If an error occurs, the function may return any value listed in the Errors section.

**Errors:**

EBUSY Attempt to join a thread which has been joined by another thread.

EDEADLK Attempt to join to yourself.

EFAULT A fault occurred when the kernel tried to access *status*.

EINTR The call was interrupted by a signal.

EINVAL Attempt to join a thread which is detached (see *ThreadDetach()*).

ESRCH The thread indicated by *tid* doesn't exist.

ETIMEDOUT A kernel timeout unblocked the call. See *TimerTimeout()*.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*pthread\_join(), pthread\_timedjoin(), ThreadCreate(), ThreadDetach()*

## Synopsis:

```
#include <time.h>

time_t time(time_t * tloc);
```

## Arguments:

*tloc*      NULL, or a pointer to a **time\_t** object where the function can store the current calendar time.

## Library:

**libc**

## Description:

The *time()* function takes a pointer to **time\_t** as an argument and returns a value of **time\_t** on exit. The returned value is the current calendar time, in seconds, since the Unix Epoch, 00:00:00 January 1, 1970 Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time (GMT)).

You typically use the **date** command to set the computer's internal clock using Coordinated Universal Time (UTC). Use the **TZ** environment variable or **\_CS\_TIMEZONE** configuration string to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

## Returns:

The current calendar time, in seconds, since 00:00:00 January 1, 1970 Coordinated Universal Time (UTC). If *tloc* isn't NULL, the current calendar time is also stored in the object pointed to by *tloc*.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

int main(void)
{
 time_t time_of_day;

 time_of_day = time(NULL);
 printf("It is now: %s", ctime(&time_of_day));
 return EXIT_SUCCESS;
}
```

produces the output:

It is now: Wed Jun 30 09:09:33 1999

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*asctime(), asctime\_r(), clock(), clock\_gettime(), ctime(), difftime(), gmtime(), localtime(), localtime\_r(), mktime(), strftime(), tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*



**Synopsis:**

```
#include <signal.h>
#include <time.h>

int timer_create(clockid_t clock_id,
 struct sigevent * evp,
 timer_t * timerid);
```

**Arguments:**

*clock\_id*     The clock source that you want to use; one of:

- `CLOCK_REALTIME` — the standard POSIX-defined timer.
- `CLOCK_SOFTTIME` — currently, the same as `CLOCK_REALTIME`.

*evp*            NULL, or a pointer to a `sigevent` structure containing the event that you want to deliver when the timer fires.

*timerid*       A pointer to a `timer_t` object where the function stores the ID of the new timer.

**Library:**

`libc`

**Description:**

The `timer_create()` function creates a per-process timer using the specified clock source, *clock\_id*, as the timing base.

You can use the time ID that the function stores in *timerid* in subsequent calls to `timer_gettime()`, `timer_settime()`, and `timer_delete()`.

The timer is created in the disabled state, and isn't enabled until you call `timer_settime()`.

We recommend the following event types:

- SIGEV\_SIGNAL
- SIGEV\_SIGNAL\_CODE
- SIGEV\_SIGNAL\_THREAD
- SIGEV\_PULSE

If the *evp* argument is NULL, a SIGALRM signal is sent to your process when the timer expires. To specify a handler for this signal, call *sigaction()*.

## Returns:

- 0 Success. The *timerid* argument is set to the timer's ID.
- 1 An error occurred (*errno* is set).

## Errors:

- EAGAIN All timers are in use. You'll have to wait for a process to release one.
- EINVAL The *clock\_id* isn't one of the valid CLOCK\_\* constants.

## Examples:

```
/*
 * Demonstrate how to set up a timer that, on expiry,
 * sends us a pulse. This example sets the first
 * expiry to 1.5 seconds and the repetition interval
 * to 1.5 seconds.
 */

#include <stdio.h>
#include <time.h>
#include <sys/netmgr.h>
#include <sys/neutrino.h>

#define MY_PULSE_CODE _PULSE_CODE_MINAVAIL

typedef union {
 struct _pulse pulse;
 /* your other message structures would go
 here too */
}
```

```
 } my_message_t;

main()
{
 struct sigevent event;
 struct itimerspec itime;
 timer_t timer_id;
 int chid;
 int rcvid;
 my_message_t msg;

 chid = ChannelCreate(0);

 event.sigev_notify = SIGEV_PULSE;
 event.sigev_coid = ConnectAttach(ND_LOCAL_NODE, 0,
 chid,
 _NTO_SIDE_CHANNEL, 0);
 event.sigev_priority = getprio(0);
 event.sigev_code = MY_PULSE_CODE;
 timer_create(CLOCK_REALTIME, &event, &timer_id);

 itime.it_value.tv_sec = 1;
 /* 500 million nsecs = .5 secs */
 itime.it_value.tv_nsec = 500000000;
 itime.it_interval.tv_sec = 1;
 /* 500 million nsecs = .5 secs */
 itime.it_interval.tv_nsec = 500000000;
 timer_settime(timer_id, 0, &itime, NULL);

 /*
 * As of the timer_settime(), we will receive our pulse
 * in 1.5 seconds (the itime.it_value) and every 1.5
 * seconds thereafter (the itime.it_interval)
 */

 for (;;) {
 rcvid = MsgReceive(chid, &msg, sizeof(msg), NULL);
 if (rcvid == 0) { /* we got a pulse */
 if (msg.pulse.code == MY_PULSE_CODE) {
 printf("we got a pulse from our timer\n");
 } /* else other pulses ... */
 } /* else other messages ... */
 }
}
```

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The QNX Neutrino version of *timer\_create()* is different from the QNX 4 version, which was based on a draft standard.

**See also:**

*clock\_getres()*, *clock\_gettime()*, *clock\_settime()*, *nanosleep()*, **\_pulse**, *sigaction()*, **sigevent**, *sleep()*, *TimerCreate()*, *timer\_delete()*, *timer\_getexpstatus()*, *timer\_getoverrun()*, *timer\_gettime()*, *timer\_settime()*

## Synopsis:

```
#include <time.h>

int timer_delete(timer_t timerid);
```

## Arguments:

*timerid*     A `timer_t` object that holds a timer ID, as set by `timer_create()`.

## Library:

`libc`

## Description:

The `timer_delete()` function removes a previously attached timer based upon the *timerid* returned from the `timer_create()` function. The timer is removed from the active system timer list, and returned to the free list of available timers.

## Returns:

0     Success.

-1     An error occurred (*errno* is set).

## Errors:

EINVAL     The timer *timerid* isn't attached to the calling process.

## Classification:

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*clock\_getres(), clock\_gettime(), clock\_settime(), nanosleep(), sleep(), timer\_create(), timer\_getexpstatus(), timer\_getoverrun(), timer\_gettime(), timer\_settime()*

**Synopsis:**

```
#include <time.h>

int timer_getexpstatus(timer_t timerid);
```

**Arguments:**

*timerid*     A `timer_t` object that holds a timer ID, as set by `timer_create()`.

**Library:**

`libc`

**Description:**

The `timer_getexpstatus()` function gets the expiry status of the time with the ID given by *timerid*.

**Returns:**

0     The timer has expired.  
-1     An error occurred (*errno* is set).

**Errors:**

EINVAL     The timer specified by *timerid* doesn't exist.

**Classification:**

POSIX 1003.1j (draft)

**Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*,  
*timer\_settime()*, *TimerInfo()*



## Synopsis:

```
#include <signal.h>
#include <time.h>

int timer_getoverrun(timer_t timerid);
```

## Arguments:

*timerid*     A `timer_t` object that holds a timer ID, as set by `timer_create()`.

## Library:

`libc`

## Description:

When a timer expiration signal is received by a process, the `timer_getoverrun()` function returns the timer expiration overrun count for the timer specified by *timerid*.

Only a single signal is queued to the process for a given timer at any point in time. When a timer that has a signal pending expires, no signal is queued and a timer overrun occurs.

The overrun count returned is the number of extra timer expirations that occurred between the time the expiration signal was queued and when it was delivered or accepted, up to but not including `DELAYTIMER_MAX`. If the number of overruns is greater than or equal to `DELAYTIMER_MAX`, the overrun count is set to `DELAYTIMER_MAX`.

The value returned by `timer_getoverrun()` applies to the most recent expiration signal for the specified timer. If no expiration signal has been delivered, the overrun count is 0.

**Returns:**

The number of overruns, or -1 if an error occurs (*errno* is set).

**Errors:**

EINVAL     Invalid timer *timerid*.

**Classification:**

POSIX 1003.1 (Realtime Extensions)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*timer\_create()*, *timer\_delete()*, *timer\_getexpstatus()*, *timer\_gettime()*,  
*timer\_settime()*, *TimerInfo()*

### **Synopsis:**

```
#include <time.h>

int timer_gettime(timer_t timerid,
 struct itimerspec *value);
```

### **Arguments:**

*timerid*    A `timer_t` object that holds a timer ID, as set by `timer_create()`.

*value*     A pointer to a `itimerspec` structure that the function fills in with the timer's time until expiry. The structure contains at least the following members:

```
struct timespec it_value
```

A `timespec` structure that contains the amount of time left before the timer expires, or zero if the timer is disarmed. This value is expressed as the relative interval until expiration, even if the timer was armed with an absolute time.

```
struct timespec it_interval
```

A `timespec` structure that contains the timer's reload value. If nonzero, it indicates a repetitive timer period.

### **Library:**

`libc`

### **Description:**

The `timer_gettime()` function gets the amount of time left before the specified timer is to expire, along with the timer's reload value, and stores it in the space provided by the *value* argument.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EINVAL The timer *timerid* isn't attached to the calling process.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*clock\_getres()*, *clock\_gettime()*, *clock\_settime()*, *nanosleep()*, *sleep()*,  
*timer\_create()*, *timer\_delete()*, *timer\_getexpstatus()*,  
*timer\_getoverrun()*, *timer\_settime()*, **timespec**

**Synopsis:**

```
#include <time.h>

int timer_settime(timer_t timerid,
 int flags,
 struct itimerspec * value,
 struct itimerspec * ovalue);
```

**Arguments:**

- timerid* A `timer_t` object that holds a timer ID, as set by `timer_create()`.
- flags* The type of timer to arm if you aren't disarming the timer. The valid bits include:
- `TIMER_ABSTIME` — the *it\_value* represents an absolute expiration date in seconds and nanoseconds from 1970. If the date specified has already passed, the function succeeds and the expiration notice is made. If you don't set this bit, the *it\_value* represents a relative expiration period that's offset from the current system time by the specified number of seconds and nanoseconds.
- value* A pointer to a `itimerspec` structure that specifies the value that you want to set for the timer's time until expiry. For more information, see `timer_gettime()`.
- ovalue* NULL, or a pointer to a `itimerspec` structure that the function fills in with the timer's former time until expiry.

**Library:**

`libc`

## Description:

The *timer\_settime()* function sets the expiration time of the timer specified by *timerid* from the *it\_value* member of the *value* argument. If the *it\_value* structure member of *value* is zero, then the timer is disarmed.

If the *it\_interval* member of *value* is nonzero, then it specifies a repeat rate that is added to the timer once the *it\_value* period has expired. Subsequently, the timer is automatically rearmed, causing it to become continuous with a period of *it\_interval*.

If the *ovalue* parameter isn't NULL, then on return from this function it contains a value representing the previous amount of time left before the timer was to have expired, or zero if the timer was disarmed. The previous interval timer period is also stored in the *it\_interval* member.

The *timerid* is local to the calling process, and must have been created using *timer\_create()*.

## Returns:

- 0 Success.
- 1 An error occurred (*errno* is set).

## Errors:

- EFAULT A fault occurred trying to access the buffers provided.
- EINVAL The timer *timerid* isn't attached to the calling process or the number of nanoseconds specified by the *tv\_nsec* member of one of the **timespec** structures in the **itimerspec** structure pointed to by *value* is less than zero or greater than or equal to 1000 million.

## Examples:

See *timer\_create()*.

## Classification:

POSIX 1003.1 (Realtime Extensions)

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*clock\_getres()*, *clock\_gettime()*, *clock\_settime()*, *errno*, *nanosleep()*,  
*sleep()*, *timer\_create()*, *timer\_delete()*, *timer\_getexpstatus()*,  
*timer\_getoverrun()*, *timer\_gettime()*

## ***timer\_timeout()*, *timer\_timeout\_r()*** © 2004, QNX Software Systems Ltd.

*Set a timeout on a blocking state*

### **Synopsis:**

```
#include <time.h>

extern int timer_timeout(
 clockid_t id,
 int flags,
 const struct sigevent* notify,
 const struct timespec* ntime,
 struct timespec* otime);

extern int timer_timeout_r(
 clockid_t id,
 int flags,
 const struct sigevent* notify,
 const struct timespec* ntime,
 struct timespec* otime);
```

### **Arguments:**

*id* The type of timer used to implement the timeout. The possible clock types of *id* are:

CLOCK\_MONOTONIC  
A clock that always increases at a constant rate and can't be adjusted.

CLOCK\_SOFTTIME  
Same as CLOCK\_REALTIME, but if the CPU is in powerdown mode, the clock stops running.

CLOCK\_REALTIME  
A clock that maintains the system time.

*flags* A bitmask that specifies which states you want the timeout to apply to; see below.

*notify* A pointer to a **sigevent** structure that defines the event that the kernel acts on if the timeout expires; see below.



- ntime* A pointer to a `timespec` structure that specifies the timeout.
- otime* A pointer to a `timespec` structure where the function can store the actual timeout.

## Library:

`libc`

## Description:

The *timer\_timeout()* and *timer\_timeout\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

The *timer\_timeout()* function sets the timeout *ntime* on any kernel blocking state. The actual timeout that occurred is returned in *otime*. The difference between the closely related *timer\_timeout()* and *TimerTimeout()* functions is the unit of time. The time in *TimerTimeout()*'s *ntime* and *otime* arguments is in nanoseconds. When *ntime* is passed to *TimerTimeout()*, the time (in `timespec`) is converted from seconds and nanoseconds into nanoseconds. When *otime* is returned to *timer\_timeout()*, the time is converted from nanoseconds into seconds and nanoseconds.

The kernel blocking states are entered as a result of the following kernel calls:

| <b>Kernel function call</b> | <b>Blocking state</b>     |
|-----------------------------|---------------------------|
| <i>InterruptWait()</i>      | STATE_INTR                |
| <i>MsgReceivev()</i>        | STATE_RECEIVE             |
| <i>MsgSendv()</i>           | STATE_SEND or STATE_REPLY |
| <i>SignalSuspend()</i>      | STATE_SIGSUSPEND          |

*continued...*

| <b>Kernel function call</b> | <b>Blocking state</b> |
|-----------------------------|-----------------------|
| <i>SignalWaitinfo()</i>     | STATE_SIGWAITINFO     |
| <i>SyncCondvarWait()</i>    | STATE_CONDVAR         |
| <i>SyncMutexLock()</i>      | STATE_MUTEX           |
| <i>SyncSemWait()</i>        | STATE_SEM             |
| <i>ThreadJoin()</i>         | STATE_JOIN            |

The user specifies which states the timeout should apply to via a bitmask passed in the *flags* argument. The bits are defined by the following constants:

| <b>Constant</b>                       | <b>Meaning</b>                |
|---------------------------------------|-------------------------------|
| <code>_NTO_TIMEOUT_CONDVAR</code>     | Timeout on STATE_CONDVAR.     |
| <code>_NTO_TIMEOUT_JOIN</code>        | Timeout on STATE_JOIN.        |
| <code>_NTO_TIMEOUT_INTR</code>        | Timeout on STATE_INTR.        |
| <code>_NTO_TIMEOUT_MUTEX</code>       | Timeout on STATE_MUTEX.       |
| <code>_NTO_TIMEOUT_RECEIVE</code>     | Timeout on STATE_RECEIVE.     |
| <code>_NTO_TIMEOUT_REPLY</code>       | Timeout on STATE_REPLY.       |
| <code>_NTO_TIMEOUT_SEM</code>         | Timeout on STATE_SEM.         |
| <code>_NTO_TIMEOUT_SEND</code>        | Timeout on STATE_SEND.        |
| <code>_NTO_TIMEOUT_SIGSUSPEND</code>  | Timeout on STATE_SIGSUSPEND.  |
| <code>_NTO_TIMEOUT_SIGWAITINFO</code> | Timeout on STATE_SIGWAITINFO. |

For example, to set a timeout on *MsgSendv()*, specify:

```
_NTO_TIMEOUT_SEND | _NTO_TIMEOUT_REPLY
```

Once a timeout is specified using *timer\_timeout()*, it's armed and released under the following conditions:

- Armed        The kernel attempts to enter a blocking state specified in *flags*.
- Released    One of the above kernel calls completed without blocking, *or* the kernel call blocks but unblocks before the timeout expires, *or* the timeout expires.

The *timer\_timeout()* function always operates on a one-shot basis. When one of the above kernel calls returns (or is interrupted by a signal), the timeout request is removed from the system. Only one timeout per thread may be in effect at a time. A second call to *timer\_timeout()*, without calling one of the above kernel functions, replaces the existing timeout on that thread. A call with *flags* set to zero ensures that a timeout won't occur on any state. This is the default when a thread is created.

Always call *timer\_timeout()* just before the function that you wish to timeout. For example:

```
...
event.sigev_notify = SIGEV_UNBLOCK;

timeout.tv_sec = 10;
timeout.tv_nsec = 0;

timer_timeout(CLOCK_REALTIME,
 _NTO_TIMEOUT_SEND | _NTO_TIMEOUT_REPLY,
 &event, &timeout, NULL);
MsgSendv(coid, NULL, 0, NULL, 0);
...
```

If the signal handler is called between the calls to *timer\_timeout()* and *MsgSendv()*, the *timer\_timeout()* values are saved during the signal handler and then are restored when the signal handler exits.

If the timeout expires, the kernel acts upon the event specified by the **sigevent** structure pointed to by the *notify* argument. We recommend the following event types in this case:

- SIGEV\_SIGNAL
- SIGEV\_SIGNAL\_CODE
- SIGEV\_SIGNAL\_THREAD
- SIGEV\_PULSE
- SIGEV\_UNBLOCK
- SIGEV\_INTR

Only SIGEV\_UNBLOCK guarantees that the kernel call unblocks. A signal may be ignored, blocked, or accepted by another thread and a pulse can only unblock a *MsgReceivev()*. If a NULL is passed for *event* then SIGEV\_UNBLOCK is assumed. In this case, a timed out kernel call will return failure with an error of ETIMEDOUT.



---

*MsgSendv()* won't unblock on SIGEV\_UNBLOCK if the server has already received the message via *MsgReceivev()* and has specified *\_NTO\_CHF\_UNBLOCK* in the *flags* argument to its *ChannelCreate()* call. In this case, it's up to the server to do a *MsgReplyv()* or *MsgError()*.

---

The timeout:

- Is specified by the *ntime* argument.
- Is relative to the current time (when *timer\_timeout()* is called), unless *flags* includes *TIMER\_ABSTIME*, which makes the timeout occur at the absolute time set in *ntime*.
- Occurs on a clock tick (see *ClockPeriod()*) so the actual wakeup time is a minimum of:  
$$(tv\_sec \times 1000000000 + tv\_nsec) \div (\text{size of timer tick}) \text{ nanoseconds}$$
where *tv\_sec* and *tv\_nsec* are fields of the **timespec** structure (defined in **<time.h>**).

If you specify a resolution that amounts to 1.7 timer ticks, you may wakeup anywhere from 1 to 1.99999999... timer ticks.

If you don't wish to block for any time, you can pass a NULL for *ntime* in which case no timer is used, the event is assumed to be SIGEV\_UNBLOCK and an attempt to enter a blocking state as set by *flags* will immediately return with ETIMEDOUT. Although a questionable practice, this can be used to poll potential blocking kernel calls. For example, you can poll for messages using *MsgReceivev()* with an immediate timeout. A much better approach is to use multiple threads and have one block waiting for messages.

If *flags* is set to \_NTO\_TIMEOUT\_NANOSLEEP, then these calls block in the STATE\_NANOSLEEP state until the timeout (or a signal which unblocks the thread) occurs. This can be used to implement an efficient kernel sleep as follows:

```
timer_timeout(CLOCK_REALTIME, _NTO_TIMEOUT_NANOSLEEP,
 NULL, &ntime, &otime);
```

If *otime* isn't NULL and the sleep is unblocked by a signal then it contains the time remaining in the sleep.

## Blocking states

The kernel calls don't block unless \_NTO\_TIMEOUT\_NANOSLEEP is specified in *flags*. In this case, the calls block as follows:

STATE\_NANOSLEEP

The calling thread blocks for the requested time period.

## Returns:

*timer\_timeout()*

The previous flags. If an error occurs, the function returns -1 and sets *errno*.

*timer\_timeout\_r()*

The previous flags. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

|        |                                                                                                   |
|--------|---------------------------------------------------------------------------------------------------|
| EAGAIN | All kernel timer entries are in use.                                                              |
| EFAULT | A fault occurred when the kernel tried to access <i>ntime</i> , <i>otime</i> , or <i>notify</i> . |
| EINTR  | The call was interrupted by a signal.                                                             |
| EINVAL | The clock type <i>id</i> isn't one of CLOCK_MONOTONIC, CLOCK_SOFTTIME, or CLOCK_REALTIME.         |

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The timeout value *starts timing out* when *timer\_timeout()* is called, *not* when the blocking state is entered. It might be possible to get preempted after calling *timer\_timeout()* but before the blocking kernel call.

## See also:

**sigevent**, *TimerCreate()*, *TimerInfo()*, *TimerTimeout()*

**Synopsis:**

```
#include <sys/neutrino.h>

int TimerAlarm(clockid_t id,
 const struct _itimer * itime,
 struct _itimer * otime);

int TimerAlarm_r(clockid_t id,
 const struct _itimer * itime,
 struct _itimer * otime);
```

**Arguments:**

- id*           The timer type to use to implement the alarm; one of:
- **CLOCK\_REALTIME** — This is the standard POSIX-defined clock. Timers based on this clock should wake up the processor if it's in a power-saving mode.
  - **CLOCK\_SOFTTIME** — This clock is only active when the processor is *not* in a power-saving mode. For example, an application using a **CLOCK\_SOFTTIME** timer to sleep wouldn't wake up the processor when the application was due to wake up. This will allow the processor to enter a power-saving mode. While the processor isn't in a power-saving mode, **CLOCK\_SOFTTIME** behaves the same as **CLOCK\_REALTIME**.
  - **CLOCK\_MONOTONIC** — This clock always increases at a constant rate and can't be adjusted.
- itime*        NULL, or a pointer to a **\_itimer** structure that specifies the length of time to wait.
- otime*        NULL, or a pointer to a **\_itimer** structure where the function can store the old timer trigger time.

## Library:

libc

## Description:

These kernel calls set an alarm signal (SIGALRM) to be delivered to the thread waiting on the timer at the time specified by *itime*. If *otime* isn't NULL, the old timer trigger time is returned.

The *TimerAlarm()* and *TimerAlarm\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

Alarm requests aren't stacked; only a single SIGALRM may be outstanding on a timer at one time. If you call *TimerAlarm()* while an alarm is outstanding, the alarm is reset to the new value passed in *itime*.

If *itime* is NULL, any previous alarm request is canceled, and no new alarm is set.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

*TimerAlarm()* If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*TimerAlarm\_r()* EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

EAGAIN All kernel timer entries are in use.

EINVAL Invalid timer value *id*.



## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

The *alarm()*, *TimerAlarm()*, and *ualarm()* requests aren't stacked; only a single SIGALRM generator can be scheduled with these functions. If the SIGALRM signal hasn't been generated, the next call to *alarm()*, *TimerAlarm()*, or *ualarm()* reschedules it.

## See also:

*alarm()*, *TimerCreate()*, *ualarm()*

## **TimerCreate(), TimerCreate\_r()**

© 2004, QNX Software Systems Ltd.

*Create a timer for a process*

### **Synopsis:**

```
#include <sys/neutrino.h>

int TimerCreate(clockid_t id,
 const struct sigevent *event);

int TimerCreate_r(clockid_t id,
 const struct sigevent *event);
```

### **Arguments:**

*id* The timing base; supported types are:

- **CLOCK\_REALTIME** — This is the standard POSIX-defined clock. Timers based on this clock should will wake up the processor if it's in a power-saving mode.
- **CLOCK\_SOFTTIME** — This clock is only active when the processor is *not* in a power-saving mode. For example, an application using a **CLOCK\_SOFTTIME** timer to sleep wouldn't wake up the processor when the application was due to wake up. This will allow the processor to enter a power-saving mode. While the processor isn't in a power-saving mode, **CLOCK\_SOFTTIME** behaves the same as **CLOCK\_REALTIME**.
- **CLOCK\_MONOTONIC** — This clock always increases at a constant rate and can't be adjusted.

*event* NULL, or a pointer to a **sigevent** structure that contains the event to deliver when the timer fires; see below.

### **Library:**

**libc**

## Description:

The *TimerCreate()* and *TimerCreate\_r()* kernel calls create a per-process timer using the clock specified by *id* as the timing base.

These functions are identical except in the way they indicate errors. See the Returns section for details.

Use the returned timer ID in subsequent calls to the other timer functions.

The timer is created in the disabled state, and isn't enabled until you call *TimerSettime()*.

The **sigevent** structure pointed to by *event* contains the event to deliver when the timer fires. We recommend the following event types in this case:

- If your process executes in a loop using *MsgReceivev()*, then SIGEV\_PULSE is a convenient way of receiving timer pulses.
- If you use signals for event notification, note that signals are always delivered to the process and not directly to the thread that created or armed the timer. You can change this by using a *sigev\_notify* of SIGEV\_SIGNAL\_THREAD.
- The notify types of SIGEV\_UNBLOCK and SIGEV\_INTR, while allowed, are of questionable use with timers. SIGEV\_UNBLOCK is typically used by the *TimerTimeout()* kernel call, and SIGEV\_INTR is typically used with the *InterruptWait()* kernel call.

If the *event* argument is NULL, a SIGALRM signal is sent to your process when the timer expires. To specify a handler for this signal, call *sigaction()*.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

*TimerCreate()* The timer ID of the newly created timer. If an error occurs, -1 is returned and *errno* is set.

*TimerCreate\_r()* The timer ID of the newly created timer. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

EINVAL The clock ID isn't valid.

EAGAIN All kernel timer objects are in use.

EFAULT A fault occurred when the kernel tried to access the buffers provided.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*sigevent*, *TimerAlarm()*, *TimerDestroy()*, *TimerInfo()*,  
*TimerSettime()*, *TimerTimeout()*

## ***TimerDestroy()*, *TimerDestroy\_r()*** © 2004, QNX Software Systems Ltd.

*Destroy a process timer*

### **Synopsis:**

```
#include <sys/neutrino.h>

int TimerDestroy(timer_t id);

int TimerDestroy_r(timer_t id);
```

### **Arguments:**

*id* The ID of the timer that you want to destroy, as returned by *TimerCreate()*.

### **Library:**

libc

### **Description:**

These kernel calls remove a previously created timer specified by *id*. The timer is removed from the active system timer list and returned to the list of available timers.

The *TimerDestroy()* and *TimerDestroy\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

If a timeout is pending when *TimerDestroy()* removes the timer, the timer is removed without being activated.

### **Blocking states**

These calls don't block.

### **Returns:**

The only difference between these functions is the way they indicate errors:

*TimerDestroy()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

*TimerDestroy\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

**Errors:**

EINVAL     The timer specified by *id* doesn't exist.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*TimerCreate()*

## ***TimerInfo(), TimerInfo\_r()***

© 2004, QNX Software Systems Ltd.

*Get information about a timer*

### **Synopsis:**

```
#include <sys/neutrino.h>

int TimerInfo(pid_t pid,
 timer_t id,
 int flags,
 struct _timer_info* info);

int TimerInfo_r(pid_t pid,
 timer_t id,
 int flags,
 struct _timer_info* info);
```

### **Arguments:**

- pid* The process ID that you're requesting the timer information for.
- id* The ID of the timer, as returned by *TimerCreate()*.
- flags* Supported flags are:
- `_NTO_TIMER_SEARCH` — if this flag is specified and the timer ID doesn't exist, return information on the *next* timer ID. This provides a mechanism to discover all of the timers in the process.
  - `_NTO_RESET_OVERRUNS` — reset the overrun count to zero in the `_timer_info` structure.
- info* A pointer to a `_timer_info` structure where the function can store the information about the specified timer. For more details, see "`struct _timer_info`," below.

### **Library:**

`libc`



**Description:**

These kernel calls get information about a previously created timer specified by *id*, and stores the information in the buffer pointed to by *info*.

The *TimerInfo()* and *TimerInfo\_r()* functions are identical except in the way they indicate errors. See the Returns section for details.

**struct \_timer\_info**

The `_timer_info` structure pointed to by *info* contains at least these members:

`uint32_t flags`

One or more of these bit flags:

`_NTO_TLACTIVE`

The timer is active.

`_NTO_TLABSOLUTE`

The timer is waiting for an absolute time to occur; otherwise, the timer is relative.

`_NTO_TLEXPARED`

The timer has expired.

`int32_t tid` The thread to which the timer is directed (0 if it's directed to the process).

`int32_t notify`

The notify type.

`clockid_t clockid`

The type of clock used.

`uint32_t overruns`

The number of overruns.

`struct sigevent event`

The event dispatched when the timer expires.

`struct itimerspec itime`

Time when the timer was started.

`struct itimerspec otime`

Time remaining before the timer expires.

For more information, see the description of *TimerCreate()*.

## Blocking states

These calls don't block.

## Returns:

The only difference between these functions is the way they indicate errors:

*TimerInfo()* The ID of the timer that the information is for. If an error occurs, -1 is returned and *errno* is set.

*TimerInfo\_r()* The ID of the timer that the information is for. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

EINVAL The timer specified by *id* doesn't exist.

ESRCH The process specified by *pid* doesn't exist.

## Classification:

QNX Neutrino

### Safety

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | No  |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*sigevent*, *TimerCreate()*

## **TimerSettime(), TimerSettime\_r()** © 2004, QNX Software Systems Ltd.

*Set the expiration time for a timer*

### **Synopsis:**

```
#include <sys/neutrino.h>

int TimerSettime(timer_t id,
 int flags,
 const struct _itimer * itime,
 struct _itimer * oitime);

int TimerSettime_r(timer_t id,
 int flags,
 const struct _itimer * itime,
 struct _itimer * oitime);
```

### **Arguments:**

- id* The ID of the timer whose an expiration date you want to set, as returned by *TimerCreate()*.
- flags* The only supported flag is `TIMER_ABSTIME`. If specified, then *nsec* represents an “absolute” expiration date in nanoseconds from the Unix Epoch, 00:00:00 January 1, 1970 UTC. If the date specified has already passed, then the expiration event is delivered immediately.
- If the flag isn’t specified, *nsec* represents a “relative” expiration period that’s offset from the given clock’s current system time in nanoseconds.
- itime* A pointer to a `_itimer` structure that specifies the expiration date. For detailed information, see “Expiration date,” below.
- oitime* NULL, or a pointer to a `_itimer` structure where the function can store the interval timer period (i.e. previous amount of time left before the timer was to have expired), or zero if the timer was disarmed at the time of the call. The previous interval timer period is also stored in the *interval\_nsec* member.

**Library:**

`libc`

**Description:**

The *TimerSettime()* and *TimerSettime\_r()* kernel calls set the expiration time of the timer specified by *id*.

These functions are identical except in the way they indicate errors. See the Returns section for details.

**Expiration date**

The expiration is specified by the *itime* argument. The `_itimer` structure contains at least the following members:

`uint64_t nsec`

The expiration time to set.

`uint64_t interval_nsec`

The interval reload time.

If the *nsec* member of *itime* is zero, then the timer is disarmed.

If the *interval\_nsec* member of *itime* is nonzero, then it specifies a repeat rate which is added to the timer once the *nsec* period has expired. Subsequently, the timer is automatically rearmed, causing it to become repetitive with a period of *interval\_nsec*.

If the timer is already armed when you call *TimerSettime()*, this call discards the previous setting and sets a new setting.

If the event notification specified by *TimerCreate()* has a *sigev\_code* of `SL_TIMER`, then at most one event is queued. In this case, if an event is pending from a previous timer when the timer fires again, a timer overrun occurs. You can use the *TimerInfo()* kernel call to obtain the number of overruns that have occurred on this timer.

## Blocking states

This call doesn't block.

## Returns:

The only difference between these functions is the way they indicate errors:

### *TimerSettime()*

If an error occurs, -1 is returned and *errno* is set. Any other value returned indicates success.

### *TimerSettime\_r()*

EOK is returned on success. This function does **NOT** set *errno*. If an error occurs, any value in the Errors section may be returned.

## Errors:

- EINVAL The timer specified by *id* doesn't exist.
- EFAULT A fault occurred when the kernel tried to access *itime* or *oitime*.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*TimerCreate(), TimerInfo()*

## ***TimerTimeout()*, *TimerTimeout\_r()*** © 2004, QNX Software Systems Ltd.

*Set a timeout on a blocking state*

### **Synopsis:**

```
#include <sys/neutrino.h>

int TimerTimeout(clockid_t id,
 int flags,
 const struct sigevent * notify,
 const uint64_t * ntime,
 uint64_t * otime);

int TimerTimeout_r(clockid_t id,
 int flags,
 const struct sigevent * notify,
 const uint64_t * ntime,
 uint64_t * otime);
```

### **Arguments:**

- id* The type of timer to implement the timeout; one of:
- **CLOCK\_REALTIME** — This is the standard POSIX-defined clock. Timers based on this clock should will wake up the processor if it's in a power-saving mode.
  - **CLOCK\_SOFTTIME** — This clock is only active when the processor is *not* in a power-saving mode. For example, an application using a **CLOCK\_SOFTTIME** timer to sleep wouldn't wake up the processor when the application was due to wake up. This will allow the processor to enter a power-saving mode. While the processor isn't in a power-saving mode, **CLOCK\_SOFTTIME** behaves the same as **CLOCK\_REALTIME**.
  - **CLOCK\_MONOTONIC** — This clock always increases at a constant rate and can't be adjusted.
- flags* Flags that specify which states the timeout applies to. For the list and description of applicable states, see the section "Timeout states."



- notify* A pointer to a **sigevent** structure that contains the event to act on when the timeout expires. See “Event types” for the list of recommended event types.
- ntime* The timeout (in nanoseconds).
- otime* A pointer to a location where the function can store the time remaining in the sleep.

## Library:

**libc**

## Description:

The *TimerTimeout()* and *TimerTimeout\_r()* kernel calls set a timeout on any kernel blocking state.

These functions are identical except in the way they indicate errors. See the Returns section for details.

These blocking states are entered as a result of the following kernel calls:

| <b>Call</b>              | <b>Blocking state</b>     |
|--------------------------|---------------------------|
| <i>InterruptWait()</i>   | STATE_INTR                |
| <i>MsgReceivev()</i>     | STATE_RECEIVE             |
| <i>MsgSendv()</i>        | STATE_SEND or STATE_REPLY |
| <i>SignalSuspend()</i>   | STATE_SIGSUSPEND          |
| <i>SignalWaitinfo()</i>  | STATE_SIGWAITINFO         |
| <i>SyncCondvarWait()</i> | STATE_CONDVAR             |
| <i>SyncMutexLock()</i>   | STATE_MUTEX               |
| <i>SyncSemWait()</i>     | STATE_SEM                 |

*continued...*

| <b>Call</b>         | <b>Blocking state</b> |
|---------------------|-----------------------|
| <i>ThreadJoin()</i> | STATE_JOIN            |

### **Timeout states**

You can specify which states the timeout should apply to via a bitmask passed in the *flags* argument. The bits are defined by the following constants:

- `_NTO_TIMEOUT_CONDVAR`  
Timeout on STATE\_CONDVAR.
- `_NTO_TIMEOUT_JOIN`  
Timeout on STATE\_JOIN.
- `_NTO_TIMEOUT_INTR`  
Timeout on STATE\_INTR.
- `_NTO_TIMEOUT_MUTEX`  
Timeout on STATE\_MUTEX.
- `_NTO_TIMEOUT_RECEIVE`  
Timeout on STATE\_RECEIVE.
- `_NTO_TIMEOUT_REPLY`  
Timeout on STATE\_REPLY.
- `_NTO_TIMEOUT_SEM`  
Timeout on STATE\_SEM.
- `_NTO_TIMEOUT_SEND`  
Timeout on STATE\_SEND.
- `_NTO_TIMEOUT_SIGSUSPEND`  
Timeout on STATE\_SIGSUSPEND.
- `_NTO_TIMEOUT_SIGWAITINFO`  
Timeout on STATE\_SIGWAITINFO.

For example, to set a timeout on *MsgSendv()*, specify:

```
_NTO_TIMEOUT_SEND | _NTO_TIMEOUT_REPLY
```

Once a timeout is specified using *TimerTimeout()*, it's armed and released under the following conditions:

- |          |                                                                                                                                                                    |
|----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Armed    | The kernel attempts to enter a blocking state specified in <i>flags</i> .                                                                                          |
| Released | One of the above kernel calls completed without blocking, <i>or</i> the kernel call blocks but unblocks before the timeout expires, <i>or</i> the timeout expires. |

*TimerTimeout()* always operates on a one-shot basis. When one of the above kernel calls returns (or is interrupted by a signal), the timeout request is removed from the system. Only one timeout per thread may be in effect at a time. A second call to *TimerTimeout()*, without calling one of the above kernel functions, replaces the existing timeout on that thread. A call with *flags* set to zero ensures that a timeout won't occur on any state. This is the default when a thread is created.

Always call *TimerTimeout()* just before the function that you wish to timeout. For example:

```
...
event.sigev_notify = SIGEV_UNBLOCK;

timeout = 10×1000000000;

TimerTimeout(CLOCK_REALTIME,
 _NTO_TIMEOUT_SEND | _NTO_TIMEOUT_REPLY,
 &event, &timeout, NULL);
MsgSendv(coid, NULL, 0, NULL, 0);
...
```

If the signal handler is called between the calls to *TimerTimeout()* and *MsgSendv()*, the *TimerTimeout()* values are saved during the signal handler and then are restored when the signal handler exits.

## EventTypes

If the timeout expires, the kernel acts upon the event specified in the **sigevent** structure pointed to by the *notify* argument. We recommend the following event types in this case:

- SIGEV\_SIGNAL
- SIGEV\_SIGNAL\_CODE
- SIGEV\_SIGNAL\_THREAD
- SIGEV\_PULSE
- SIGEV\_UNBLOCK
- SIGEV\_INTR

Only SIGEV\_UNBLOCK guarantees that the kernel call unblocks. A signal may be ignored, blocked, or accepted by another thread, and a pulse can only unblock a *MsgReceivev()*. If you pass NULL for *event*, SIGEV\_UNBLOCK is assumed. In this case, a timed out kernel call returns failure with an error of ETIMEDOUT.



---

*MsgSendv()* doesn't unblock on SIGEV\_UNBLOCK if the server has already received the message via *MsgReceivev()* and has specified \_NTO.CHF.UNBLOCK in the *flags* argument to its *ChannelCreate()* call. In this case, it's up to the server to do a *MsgReplyv()*.

---

## The timeout

The type of timer used to implement the timeout is specified with the *id* argument.

The timeout:

- Is specified by the *ntime* argument (the number of nanoseconds).
- Is relative to the current time (when *TimerTimeout()* is called), unless *flags* includes TIMER\_ABSTIME, which makes the timeout occur at the absolute time set in *ntime*.

- Occurs on a clock tick (see *ClockPeriod()*) so the actual wakeup time is a minimum of:  
 $( ntime ) \div ( \text{size of timer tick} ) \text{ nanoseconds}$

If you specify a resolution that amounts to 1.7 timer ticks, you'll wake up in at least 1.7 timer ticks.

If you don't wish to block for any time, you can pass a NULL for *ntime*, in which case no timer is used, the event is assumed to be SIGEV\_UNBLOCK, and an attempt to enter a blocking state as set by *flags* immediately returns with ETIMEDOUT. Although a questionable practice, you can use it to poll potential blocking kernel calls. For example, you can poll for messages using *MsgReceivev()* with an immediate timeout. A much better approach is to use multiple threads and have one block waiting for messages.

If you set *flags* to \_NTO\_TIMEOUT\_NANOSLEEP, then these calls block in the STATE\_NANOSLEEP state until the timeout (or a signal that unblocks the thread) occurs. You can use this to implement an efficient kernel sleep as follows:

```
TimerTimeout(CLOCK_REALTIME, _NTO_TIMEOUT_NANOSLEEP,
 NULL, ntime, otime);
```

If *otime* isn't NULL and the sleep is unblocked by a signal, it contains the time remaining in the sleep.

## Blocking states

These calls don't block unless you specify \_NTO\_TIMEOUT\_NANOSLEEP in *flags*. In this case, the calls block as follows:

### STATE\_NANOSLEEP

The calling thread blocks for the requested time period.

## Returns:

The only difference between these functions is the way they indicate errors:

### *TimerTimeout()*

The previous flags. If an error occurs, -1 is returned and *errno* is set.

### *TimerTimeout\_r()*

The previous flags. This function does **NOT** set *errno*. If an error occurs, the negative of a value from the Errors section is returned.

## Errors:

|        |                                                                                                   |
|--------|---------------------------------------------------------------------------------------------------|
| EAGAIN | All kernel timer entries are in use.                                                              |
| EFAULT | A fault occurred when the kernel tried to access <i>ntime</i> , <i>otime</i> , or <i>notify</i> . |
| EINTR  | The call was interrupted by a signal.                                                             |
| EINVAL | Invalid timer value <i>id</i> .                                                                   |

## Classification:

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

`sigevent`, *TimerCreate()*, *TimerInfo()*

# ***times()***

© 2004, QNX Software Systems Ltd.

*Get time-accounting information*

## **Synopsis:**

```
#include <sys/times.h>

clock_t times(struct tms* buffer);
```

## **Arguments:**

*buffer* A pointer to a **tms** structure where the function can store the time-accounting information. For information about the **tms** structure, see below.

## **Library:**

**libc**

## **Description:**

The *times()* function stores time-accounting information in the structure pointed to by *buffer*. The type **clock\_t** and the **tms** structure are defined in the **<sys/times.h>** header file.

The **tms** structure contains at least the following members:

**clock\_t** *tms\_utime*

The CPU time charged for the execution of user instructions of the calling process.

**clock\_t** *tms\_stime*

The CPU time charged for execution by the system on behalf of the calling process.

**clock\_t** *tms\_cutime*

The sum of the *tms\_utime* and *tms\_cutime* values of the child processes.

**clock\_t** *tms\_cstime*

The sum of the *tms\_stime* and *tms\_cstime* values of the child processes.



All times are in CLK\_TCK'ths of a second. CLK\_TCK is defined in the <time.h> header file. A CLK\_TCK is the equivalent of:

```
#define sysconf(_SC_CLK_TCK)
```

The times of a terminated child process are included in the *tms\_cutime* and *tms\_cstime* elements of the parent when a *wait()* or *waitpid()* function returns the process ID of this terminated child. If a child process hasn't waited for its terminated children, their times aren't included in its times.

## Returns:

The elapsed real time, in clock ticks, of kernel uptime.




---

The value returned may overflow the possible range of type `clock_t`.

---

## Examples:

```
/*
 * The following program executes the program
 * specified by argv[1]. After the child program
 * is finished, the cpu statistics of the child are
 * printed.
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/times.h>

int main(int argc, char **argv)
{
 struct tms chldtim;

 system(argv[1]);
 times(&chldtim);
 printf("system time = %d\n", chldtim.tms_cstime);
 printf("user time = %d\n", chldtim.tms_cutime);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*clock\_gettime()*

## Synopsis:

```
#include <time.h>

struct timespec {
 time_t tv_sec;
 long tv_nsec;
}
```

## Description:

The **timespec** structure specifies a time in seconds and nanoseconds. The members include:

*tv\_sec*      The number of seconds. If specifying an absolute time, this member is the number of seconds since 1970.

*tv\_nsec*     The number of nanoseconds.

## Classification:

POSIX 1003.1 (Realtime Extensions)

## See also:

*nsec2timespec()*, *timespec2nsec()*

## ***timespec2nsec()***

© 2004, QNX Software Systems Ltd.

*Convert a `timespec` structure to nanoseconds*

### **Synopsis:**

```
#include <time.h>

_uint64 timespec2nsec(const struct timespec* ts);
```

### **Arguments:**

*ts* A pointer to the `timespec` that you want to convert to nanoseconds.

### **Library:**

`libc`

### **Description:**

The *timespec2nsec()* function converts the number of seconds and nanoseconds in the `timespec` structure pointed to by *ts* into nanoseconds.

### **Returns:**

The number of nanoseconds.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*nsec2timespec()*, **timespec**

## *timezone*

© 2004, QNX Software Systems Ltd.

*The number of seconds by which the local time zone is earlier than UTC*

### **Synopsis:**

```
#include <time.h>

long int timezone;
```

### **Description:**

This global variable holds the number of seconds by which the local time zone is earlier than Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time). Whenever you call a time function, *tzset()* is called to set the variable, based on the current time zone.

### **Classification:**

QNX Neutrino

### **See also:**

*daylight, tzname, tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the Neutrino *User’s Guide*

**Synopsis:**

```
#include <time.h>
struct tm {
 int tm_sec;
 int tm_min;
 int tm_hour;
 int tm_mday;
 int tm_mon;
 int tm_year;
 int tm_wday;
 int tm_yday;
 int tm_isdst;
 long int tm_gmtoff;
 const char * tm_zone;
};
```

**Description:**

The **tm** structure describes the calendar time. The members of this structure include:

|                 |                                                                           |
|-----------------|---------------------------------------------------------------------------|
| <i>tm_sec</i>   | Seconds after the minute, in the range [0,61], allowing for leap seconds. |
| <i>tm_min</i>   | Minutes after the hour, in the range [0,59].                              |
| <i>tm_hour</i>  | Hours after midnight, in the range [0,23].                                |
| <i>tm_mday</i>  | Day of the month, in the range [1,31].                                    |
| <i>tm_mon</i>   | Months since January, in the range [0,11].                                |
| <i>tm_year</i>  | Years since 1900.                                                         |
| <i>tm_wday</i>  | Days since Sunday, in the range [0,6].                                    |
| <i>tm_yday</i>  | Days since January 1, in the range [0,365], allowing for leap years.      |
| <i>tm_isdst</i> | Daylight saving time flag.                                                |

*tm\_gmtoff*      Offset from UTC — see *setlocale()*.

*tm\_zone*        String for the time zone name.

### **Classification:**

ANSI

### **See also:**

*asctime()*, *gmtime()*, *gmtime\_r()*, *localtime()*, *localtime\_r()*, *mktime()*,  
*setlocale()*, *strftime()*, *wcsftime()*



## Synopsis:

```
#include <stdio.h>

FILE* tmpfile(void);

FILE* tmpfile64(void);
```

## Library:

libc

## Description:

The *tmpfile()* and *tmpfile64()* functions create a temporary file and opens a corresponding **FILE** stream. The file is automatically removed when it's closed or when the program terminates. The file is opened in update mode (as in *fopen()*'s **w+** mode).

If the process is killed between file creation and unlinking, a permanent file may be left behind.



When a stream is opened in update mode, both reading and writing may be performed. However, writing may not be followed by reading without an intervening call to the *fflush()* function, or to a file-positioning function (*fseek()*, *fsetpos()*, *rewind()*). Similarly, reading may not be followed by writing without an intervening call to a file-positioning function, unless the read resulted in end-of-file.

---

## Returns:

A pointer to the stream of the temporary file, or NULL if an error occurs (*errno* is set).

## Errors:

**EACCESS**      The calling process doesn't have permission to create the temporary file.

- EMFILE      The calling process already has already used OPEN\_MAX file descriptors.
- ENFILE      The system already has the maximum number of files open.
- EROFS      The filesystem for the temporary file is read-only.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

static FILE *TempFile;

int main(void)
{
 TempFile = tmpfile();
 ...
 fclose(TempFile);

 /* The temporary file will be removed when we exit. */
 return EXIT_SUCCESS;
}
```

**Classification:**

*tmpfile()* is ANSI, *tmpfile64()* is for large-file support

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fopen(), fopen64(), freopen(), freopen64(), tempnam(), tmpnam()*

## ***tmpnam()***

© 2004, QNX Software Systems Ltd.

*Generate a unique string for use as a filename*

### **Synopsis:**

```
#include <stdio.h>

char* tmpnam(char* buffer);
```

### **Arguments:**

*buffer* NULL, or a pointer to a buffer where the function can store the filename. If *buffer* isn't NULL, the buffer must be at least *L\_tmpnam* bytes long.

### **Library:**

libc

### **Description:**

The *tmpnam()* function generates a unique string that's a valid filename and that's not the same as the name of an existing file.

The *tmpnam()* function generates up to TMP\_MAX unique file names before it starts to recycle them.

The generated filename is prefixed with the first accessible directory contained in:

- The **TMPDIR** environment variable
- The temporary file directory *P\_tmpdir* (defined in `<stdio.h>`)
- The `_PATH_TMP` constant (defined in `<paths.h>`)

If all of these paths are inaccessible, *tmpnam()* attempts to use `/tmp` and then the current working directory.

The generated filename is stored in an internal buffer; if *buffer* is NULL, the function returns a pointer to this buffer; otherwise, *tmpnam()* copies the filename into *buffer*.

Subsequent calls to *tmpnam()* reuse the internal buffer. If *buffer* is NULL, you might want to duplicate the resulting string. For example,

```
char *name1, *name2;

name1 = strdup(tmpnam(NULL));
name2 = strdup(tmpnam(NULL));
```

## Returns:

A pointer to the generated filename for success, or NULL if an error occurs (*errno* is set).

## Examples:

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 char filename[L_tmpnam];
 FILE *fp;

 tmpnam(filename);
 fp = fopen(filename, "w+b");
 ...
 fclose(fp);
 remove(filename);

 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | Yes                     |
| Interrupt handler  | No                      |
| Signal handler     | No                      |
| Thread             | Read the <i>Caveats</i> |

## **Caveats:**

The *tmpnam()* function *isn't thread-safe* if you pass it a NULL *buffer*.

This function only creates pathnames; the application must create and remove the files.

It's possible for another thread or process to create a file with the same name between when the pathname is created and the file is opened.

## **See also:**

*tempnam()*, *tmpfile()*

## Synopsis:

```
#include <ctype.h>

int tolower(int c);
```

## Arguments:

*c* The character that you want to convert.

## Library:

libc

## Description:

The *tolower()* function converts *c* to a lowercase letter, if *c* represents an uppercase letter.

## Returns:

The corresponding lowercase letter when the argument is an uppercase letter; otherwise, the original character is returned.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char chars[] = {
 'A',
 '5',
 '$',
 'Z'
};

#define SIZE sizeof(chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
```

```
printf("%c ", tolower(chars[i]));
}
printf("\n");
return EXIT_SUCCESS;
}
```

produces the output:

```
a 5 $ z
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *strlwr()*, *strupr()*, *toupper()*



## Synopsis:

```
#include <ctype.h>

int toupper(int c);
```

## Arguments:

*c* The character that you want to convert.

## Library:

libc

## Description:

The *toupper()* function converts *c* to a uppercase letter, if *c* represents a lowercase letter.

## Returns:

The corresponding uppercase letter when the argument is a lowercase letter; otherwise, the original character is returned.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

char chars[] = {
 'a',
 '5',
 '$',
 'z'
};

#define SIZE sizeof(chars) / sizeof(char)

int main(void)
{
 int i;

 for(i = 0; i < SIZE; i++) {
```

```
printf("%c ", toupper(chars[i]));
}
printf("\n");
return EXIT_SUCCESS;
}
```

produces the output:

A 5 \$ Z

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

*isalnum()*, *isalpha()*, *iscntrl()*, *isdigit()*, *isgraph()*, *islower()*, *isprint()*, *ispunct()*, *isspace()*, *isupper()*, *isxdigit()*, *strlwr()*, *strupr()*, *tolower()*

## Synopsis:

```
#include <wctype.h>

wint_t towctrans(wint_t wc,
 wctrans_t category);
```

## Arguments:

*wc*            The wide character that you want to convert.

*category*      How you want to convert the character; get this by calling *wctrans()*.

## Library:

`libc`

## Description:

The *towctrans()* function converts *wc*, using the mapping described by *category*. The following functions are equivalent:

| <b>Function</b>             | <b>Equivalent <i>wctrans()</i> call</b>          |
|-----------------------------|--------------------------------------------------|
| <code>towlower( wc )</code> | <code>towctrans( wc, wctrans("tolower") )</code> |
| <code>toupper( wc )</code>  | <code>towctrans( wc, wctrans("toupper") )</code> |

## Returns:

The corresponding converted wide character when the argument is valid; otherwise, the original wide character.

## Errors:

EINVAL      The conversion descriptor in *category* is invalid.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*wctrans()*

“Character manipulation functions” and “Wide-character functions”  
in *Library Reference Summary*

**Synopsis:**

```
#include <wctype.h>

wint_t tolower(wint_t wc);
```

**Arguments:**

*wc*     The wide character that you want to convert.

**Library:**

`libc`

**Description:**

The *tolower()* function converts *wc* to a lowercase letter, if *wc* represents an uppercase letter.

**Returns:**

The corresponding lowercase letter when the argument is an uppercase letter; otherwise, the original wide character is returned.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“Character manipulation functions” and “Wide-character functions”  
in *Library Reference Summary*

**Synopsis:**

```
#include <wctype.h>

wint_t toupper(wint_t wc);
```

**Arguments:**

*wc*     The wide character that you want to convert.

**Library:**

`libc`

**Description:**

The *toupper()* function converts *wc* to an uppercase letter if *wc* represents a lowercase letter.

**Returns:**

The corresponding uppercase letter when the argument is a lowercase letter; otherwise, the original wide character is returned.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“Character manipulation functions” and “Wide-character functions”  
in *Library Reference Summary*



## Synopsis:

```
#include <sys/neutrino.h>

int TraceEvent(int mode,
 ...);
```

## Arguments:

*mode*     A command that indicates what you want to trace. Certain modes require additional arguments.

## Library:

libc

## Description:

The *TraceEvent()* function controls all stages of the instrumentation process such as initialization, starting, execution control and stopping. These stages consist of the following activities:

- creating internal circular link list of trace buffers
- initializing filters
- turning on or off the event stream
- deallocating the internal circular link list of trace buffers



---

This function requires the instrumented kernel. For more information, see the documentation for the System Analysis Toolkit (SAT).

---

## Returns:

If *mode* is set to `_NTO_TRACE_QUERYEVENTS`

Number of events in the buffer, or -1 if an error occurs (*errno* is set).

If *mode* isn't set to `_NTO_TRACE_QUERYEVENTS`

0 for success, or -1 if an error occurs (*errno* is set).

## Errors:

|           |                                                                                |
|-----------|--------------------------------------------------------------------------------|
| ECANCELED | The requested action has been canceled.                                        |
| EFAULT    | The requested action has been specified out of order.                          |
| ENOMEM    | Insufficient memory to allocate the trace buffers.                             |
| ENOSUP    | The requested action isn't supported.                                          |
| EPERM     | The application doesn't have the appropriate permission to perform the action. |

## Classification:

QNX Neutrino

### Safety

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | Read the <i>Caveats</i> |
| Signal handler     | Yes                     |
| Thread             | Yes                     |

## Caveats:

You can call *TraceEvent()* from an interrupt/event handler. However, not all trace modes are valid in this case. The valid trace modes are:

- `_NTO_TRACE_INSERTSUSEREVENT`
- `_NTO_TRACE_INSERTCUSEREVENT`
- `_NTO_TRACE_INSERTUSRSTREVENT`

- `_NTO_TRACE_INSERTEVENT`
- `_NTO_TRACE_STOP`
- `_NTO_TRACE_STARTNOSTATE`
- `_NTO_TRACE_START`

**See also:**

*InterruptAttach(), InterruptHookTrace()*

## ***truncate()***

© 2004, QNX Software Systems Ltd.

*Truncate a file to a specified length*

---

### **Synopsis:**

```
#include <unistd.h>

int truncate(const char* path,
 off_t length);
```

### **Arguments:**

*path*        The path name of the file that you want to truncate.  
*length*      The new size of the file.

### **Library:**

`libc`

### **Description:**

The *truncate()* function causes the regular file named by *path* to have a size of *length* bytes.

The effect of *truncate()* on other types of files is unspecified. If the file previously was larger than *length*, the extra data is lost. If it was previously shorter than *length*, bytes between the old and new lengths are read as zeroes. The process must have write permission for the file.

If the request would cause the file size to exceed the soft file size limit for the process, the request fails and the implementation generates the SIGXFSZ signal for the process.

This function doesn't modify the file offset for any open file descriptions associated with the file. On successful completion, if the file size is changed, *truncate()* marks for update the *st\_ctime* and *st\_mtime* fields of the file, and if the file is a regular file, the S\_ISUID and S\_ISGID bits of the file mode may be cleared.

**Returns:**

- 0 Success.
- 1 An error occurred; *errno* is set.

**Errors:**

- EACCES A component of the path prefix denies search permission, or write permission is denied on the file.
- EFAULT The *path* argument points outside the process's allocated address space.
- EFBIG The *length* argument was greater than the maximum file size.
- EINTR A signal was caught during execution.
- EINVAL The *length* argument is invalid, or the *path* argument isn't an ordinary file.
- EIO An I/O error occurred while reading from or writing to a filesystem.
- EISDIR The named file is a directory.
- ELOOP Too many symbolic links were encountered in resolving *path*.
- EMFILE The maximum number of file descriptors available to the process has been reached.
- EMULTIHOP Components of *path* require hopping to multiple remote machines and filesystem type doesn't allow it.
- ENAMETOOLONG  
The length of the specified pathname exceeds PATH\_MAX bytes, or the length of a component of the pathname exceeds NAME\_MAX bytes.

|         |                                                                                                       |
|---------|-------------------------------------------------------------------------------------------------------|
| ENFILE  | Additional space couldn't be allocated for the system file table.                                     |
| ENOENT  | A component of <i>path</i> doesn't name an existing file or <i>path</i> is an empty string.           |
| ENOLINK | The <i>path</i> argument points to a remote machine and the link to that machine is no longer active. |
| ENOTDIR | A component of the path prefix of <i>path</i> isn't a directory.                                      |
| EROFS   | The named file resides on a read-only filesystem.                                                     |

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*chmod()*, *fcntl()*, *ftruncate()*, *open()*

## Synopsis:

```
#include <unistd.h>

char *ttyname(int fildev);
```

## Arguments:

*fildev*     A file descriptor that's associated with the file whose name you want to get.

## Library:

libc

## Description:

The *ttyname()* function returns a pointer to a static buffer that contains a fully qualified pathname associated with the file associated with *fildev*.

## Returns:

A pointer to the pathname for *fildev*, or NULL if an error occurred (*errno* is set).

## Errors:

|        |                                                                                                 |
|--------|-------------------------------------------------------------------------------------------------|
| EBADF  | The <i>fildev</i> argument is invalid.                                                          |
| ENOSYS | The <i>ttyname()</i> function isn't implemented for the filesystem specified by <i>fildev</i> . |
| ENOTTY | Not a tty.                                                                                      |

## Examples:

```
/*
 * The following program prints out the name
 * of the terminal associated with stdin.
 */
```

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(void)
{
 if(isatty(0)) {
 printf("%s\n", ttyname(0));
 } else {
 printf("\n");
 }
 return EXIT_SUCCESS;
}
```

## **Classification:**

POSIX 1003.1

### **Safety**

---

|                    |    |
|--------------------|----|
| Cancellation point | No |
| Interrupt handler  | No |
| Signal handler     | No |
| Thread             | No |

## **See also:**

*ctermid()*, *setsid()*, *ttyname\_r()*



## Synopsis:

```
#include <unistd.h>

int ttyname_r(int fildes,
 char* name,
 size_t namesize);
```

## Arguments:

|                 |                                                                               |
|-----------------|-------------------------------------------------------------------------------|
| <i>fildes</i>   | A file descriptor that's associated with the file whose name you want to get. |
| <i>name</i>     | A pointer to a buffer where the function can store the path name.             |
| <i>namesize</i> | The size of the buffer.                                                       |

## Library:

libc

## Description:

The *ttyname\_r()* function stores the null-terminated pathname of the terminal associated with the file descriptor *fildes* in the character array referenced by *name*. The array is *namesize* characters long and should have space for the name and the terminating NULL character.

## Returns:

Zero for success, or an error number.

## Errors:

|        |                                                                                                   |
|--------|---------------------------------------------------------------------------------------------------|
| EBADF  | The <i>fildes</i> argument isn't a valid file descriptor.                                         |
| ENOSYS | The <i>ttyname_r()</i> function isn't implemented for the filesystem specified by <i>fildes</i> . |
| ENOTTY | The <i>fildes</i> argument doesn't refer to a tty.                                                |

ERANGE     The value of *namesize* is smaller than the length of the string to be returned, including the terminating null character.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*ctermid(), errno, setsid(), ttyname()*

## **Synopsis:**

```
#include <time.h>

char *tzname [];
```

## **Description:**

This global variable holds the standard abbreviations for the time zone and the time zone when daylight saving time is in effect. Whenever you call a time function, *tzset()* is called to set the values in the array, based on the current time zone.

## **Classification:**

QNX Neutrino

## **See also:**

*daylight, timezone, tzset()*

“Setting the time zone” in the Configuring Your Environment chapter of the *Neutrino User’s Guide*

## ***tzset()***

© 2004, QNX Software Systems Ltd.

*Set the time according to the current time zone*

### **Synopsis:**

```
#include <time.h>

void tzset(void);
```

### **Library:**

libc

### **Description:**

The *tzset()* function sets the global variables *daylight*, *timezone* and *tzname* according to the value of the **TZ** environment variable, or to the value of the **\_CS\_TIMEZONE** configuration string if **TZ** isn't set.

The global variables have the following values after *tzset()* is executed:

|                 |                                                                                                                                                                                                                                                                                                                                                          |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>daylight</i> | Zero indicates that daylight saving time isn't supported in the locale; a nonzero value indicates that daylight saving time is supported in the locale. This variable is cleared or set after a call to the <i>tzset()</i> function, depending on whether or not a daylight saving time abbreviation is specified in the <b>TZ</b> environment variable. |
| <i>timezone</i> | The number of seconds that the local time zone is earlier than Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time (GMT)).                                                                                                                                                                                                           |
| <i>tzname</i>   | A two-element array pointing to strings giving the abbreviations for the name of the time zone when standard and daylight saving time are in effect.                                                                                                                                                                                                     |

The time that you set on the computer with the **date** command reflects Coordinated Universal Time (UTC). The environment variable **TZ** is used to establish the local time zone. For more information, see "Setting the time zone" in the Configuring Your Environment chapter of the *Neutrino User's Guide*.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <time.h>

void print_zone()
{
 char *tz;

 printf("TZ: %s\n", (tz = getenv("TZ"))
 ? tz : "default EST5EDT");
 printf(" daylight: %d\n", daylight);
 printf(" timezone: %ld\n", timezone);
 printf(" time zone names: %s %s\n",
 tzname[0], tzname[1]);
}

int main(void)
{
 print_zone();
 setenv("TZ", "PST8PDT", 1);
 tzset();
 print_zone();
 return EXIT_SUCCESS;
}
```

produces the output:

```
TZ: default EST5EDT
 daylight: 1
 timezone: 18000
 time zone names: EST EDT
TZ: PST8PDT
 daylight: 1
 timezone: 28800
 time zone names: PST PDT
```

## Classification:

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*ctime(), daylight, localtime(), localtime\_r(), mktime(), strftime()  
timezone, tzname*

“Setting the time zone” in the Configuring Your Environment chapter  
of the Neutrino *User’s Guide*

## Synopsis:

```
#include <unistd.h>

useconds_t ualarm(useconds_t usec,
 useconds_t interval);
```

## Arguments:

*usec*        The number of microseconds that you want to elapse before the first alarm occurs, or 0 to cancel any previous request for an alarm.

*interval*    The number of microseconds that you want to elapse before the subsequent alarms occur.

## Library:

`libc`

## Description:

The *ualarm()* function causes the system to send the calling process a SIGALRM signal after *usec* microseconds of real-time have elapsed. The alarm is then sent every *interval* microseconds after that.

Processor scheduling delays may cause a delay between when the signal is sent and when the process actually handles it.

If *usec* is 0, any previous *ualarm()* request is canceled.

## Returns:

0        There was no previous *ualarm()* request.

-1       An error occurred (*errno* is set).

Any other value

The number of microseconds until the next scheduled SIGALRM.

**Errors:**

EAGAIN      All timers are in use; wait for a process to release one and try again.

**Examples:**

```
#include <stdio.h>
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 useconds_t timeleft;

 printf("Set the alarm and sleep\n");
 ualarm((useconds_t)(10 * 1000 * 1000), 0);
 sleep(5); /* go to sleep for 5 seconds */

 /*
 * To get the time left before the SIGALRM is
 * to arrive, one must cancel the initial timer,
 * which returns the amount of time it had
 * remaining.
 */
 timeleft = ualarm(0, 0);
 printf("Time left before cancel, and rearm: %ld\n",
 timeleft);

 /*
 * Start a new timer that kicks us when timeleft
 * seconds have passed.
 */
 ualarm(timeleft, 0);

 /*
 * Wait until we receive the SIGALRM signal; any
 * signal kills us, though, since we don't have
 * a signal handler.
 */
 printf("Hanging around, waiting to exit\n");
 pause();

 /* You'll never get here. */
 return EXIT_SUCCESS;
}
```



**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*alarm()*, *TimerAlarm()*, and *ualarm()* requests aren't "stacked"; only a single SIGALRM generator can be scheduled with these functions. If the SIGALRM signal hasn't been generated, the next call to *alarm()*, *TimerAlarm()*, or *ualarm()* reschedules it.

Don't mix calls to *ualarm()* with *nanosleep()*, *sleep()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, *timer\_settime()*, or *usleep()*.

**See also:**

*alarm()*, *nanosleep()*, *sigaction()*, *sleep()*, *timer\_create()*, *timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, *timer\_settime()*, *TimerAlarm()*, *usleep()*

# UDP

© 2004, QNX Software Systems Ltd.

*Internet User Datagram Protocol*

## Synopsis:

```
#include <sys/socket.h>
#include <netinet/in.h>

int socket(AF_INET,
 SOCK_DGRAM,
 0);
```

## Description:

UDP is a simple, unreliable datagram protocol that's used to support the SOCK\_DGRAM abstraction for the Internet protocol family. UDP sockets are connectionless and are normally used with the *sendto()* and *recvfrom()* calls, although you can also use the *connect()* call to fix the destination for future packets (in which case you can use the *recv()* or *read()* and *send()* or *write()* system calls).

UDP address formats are identical to those used by TCP. In particular, UDP provides a port identifier in addition to the normal Internet address format. Note that the UDP port space is separate from the TCP port space; that is, a UDP port may *not* be “connected” to a TCP port. In addition, broadcast packets may be sent — assuming the underlying network supports this — by using a reserved broadcast address; this address is network-interface dependent.

You can use options at the IP transport level with UDP (see the IP protocol).

## Returns:

A descriptor referencing the socket, or -1 if an error occurs (*errno* is set).

## Errors:

EADDRINUSE    You tried to create a socket with a port that has already been allocated.

EADDRNOTAVAIL

You tried to create a socket with a network address for which no network interface exists.

EISCONN

You tried to establish a connection on a socket that already has one, or to send a datagram with the destination address specified and the socket is already connected.

ENOBUFS

The system ran out of memory for an internal data structure.

ENOTCONN

You tried to send a datagram, but no destination address was specified and the socket hasn't been connected.

**See also:**

IP protocol

*connect(), getsockopt(), read(), recv(), recvfrom(), send(), sendto(), socket(), write()*

*RFC 768*

## ***ultoa()*, *ulltoa()***

© 2004, QNX Software Systems Ltd.

*Convert an unsigned long integer into a string, using a given base*

### **Synopsis:**

```
#include <stdlib.h>

char* ultoa(unsigned long int value,
 char* buffer,
 int radix);

char* ulltoa(uint64_t value
 char* buffer,
 int radix);
```

### **Arguments:**

*value*     The value to convert into a string.

*buffer*    A buffer in which the function stores the string. The size of the buffer must be at least 33 bytes when converting values in base 2 (binary).

*radix*     The base to use when converting the number. This value must be in the range:

$2 \leq \textit{radix} \leq 36$

### **Library:**

`libc`

### **Description:**

The *ultoa()* and *ulltoa()* functions convert the unsigned binary integer *value* into the equivalent string in base *radix* notation, storing the result in the character array pointed to by *buffer*. A NUL character is appended to the result.

**Returns:**

A pointer to the result.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

void print_value(unsigned long int value)
{
 int base;
 char buffer[33];

 for(base = 2; base <= 16; base = base + 2)
 printf("%2d %s\n", base,
 ultoa(value, buffer, base));
}

int main(void)
{
 print_value((unsigned) 12765L);
 return EXIT_SUCCESS;
}
```

produces the output:

```
 2 11000111011101
 4 3013131
 6 135033
 8 30735
10 12765
12 7479
14 491b
16 31dd
```

**Classification:**

*ultoa()* is QNX 4; *ulltoa()* is Unix

**Safety**

Cancellation point No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*atoi()*, *atol()*, *itoa()*, *ltoa()*, *sscanf()*, *strtol()*, *strtoul()*, *utoa()*

**Synopsis:**

```
#include <sys/types.h>
#include <sys/stat.h>

mode_t umask(mode_t cmask);
```

**Arguments:**

*cmask* The new file-mode creation mask; that is, the permissions that you don't want set when the process creates a file. The mask is a combination of these bits:

| Owner   | Group   | Others  | Permission                                                                                                                          |
|---------|---------|---------|-------------------------------------------------------------------------------------------------------------------------------------|
| S_IRUSR | S_IRGRP | S_IROTH | Read                                                                                                                                |
| S_IRWXU | S_IRWXG | S_IRWXO | Read, write, execute/search. A bitwise inclusive OR of the other three constants.<br>(S_IRWXU is OR of IRUSR, S_IWSUR and S_IXUSR.) |
| S_IWUSR | S_IWGRP | S_IWOTH | Write                                                                                                                               |
| S_IXUSR | S_IXGRP | S_IXOTH | Execute/search                                                                                                                      |

**Library:**

`libc`

**Description:**

The *umask()* function sets the process's file-mode creation mask to *cmask*, and returns the previous value of the mask. Only the file permission bits (as defined in `<sys/stat.h>`) are used.

The file-mode creation mask for the process is used when you call *creat()*, *mkdir()*, *mkfifo()*, and *open()*, to turn off permission bits in the *mode* argument supplied. Bit positions set in *cmask* are cleared in the mode of the created file.

**Returns:**

The previous value of the file-mode creation mask.

**Examples:**

```
/*
 * Set the umask to RW for owner,group; R for other
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/stat.h>

int main(void)
{
 mode_t omask;
 mode_t nmask;

 nmask = S_IRUSR | S_IWUSR | /* owner read write */
 S_IRGRP | S_IWGRP | /* group read write */
 S_IROTH; /* other read */
 omask = umask(nmask);
 printf("Mask changed from %o to %o\n",
 omask, nmask);
 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*chmod(), creat(), mkdir(), mkfifo(), open(), stat()*

# ***umount()***

© 2004, QNX Software Systems Ltd.

*Unmount a filesystem*

## **Synopsis:**

```
#include <sys/mount.h>

int umount(const char* dir,
 int flags);
```

## **Arguments:**

- dir*      The filesystem that you want to unmount.
- flags*    Flags that control the operation. Currently, the only valid value for *flags* is:
- `_MOUNT_FORCE` — force an unmount to occur.

## **Library:**

`libc`

## **Description:**

The *umount()* function sends a request to the server to unmount the path described by *dir*.

## **Returns:**

-1 on failure.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*mount()*

Writing a Resource Manager in *Programmer's Guide*

## ***UNALIGNED\_PUT16()***

© 2004, QNX Software Systems Ltd.

*Write a misaligned 16-bit value safely*

### **Synopsis:**

```
#include <gulliver.h>

void UNALIGNED_PUT16(uint16_t *loc,
 uint16_t num);
```

### **Arguments:**

*loc*      The address where you want to write the value.

*num*      The value that you want to write.

### **Library:**

`libc`

### **Description:**

The *UNALIGNED\_PUT16()* macro lets you write the value *num* at the misaligned address *loc* without faulting.

### **Classification:**

QNX Neutrino

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*UNALIGNED\_PUT16()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16(), ENDIAN\_BE32(), ENDIAN\_BE64(),  
ENDIAN\_LE16(), ENDIAN\_LE32(), ENDIAN\_LE64(),  
ENDIAN\_RET32(), ENDIAN\_RET64(), ENDIAN\_SWAP16(),  
ENDIAN\_SWAP32(), ENDIAN\_SWAP64(), htonl(), htons(), ntohl(),  
ntohs(), UNALIGNED\_PUT32(), UNALIGNED\_PUT64(),  
UNALIGNED\_RET16(), UNALIGNED\_RET32(),  
UNALIGNED\_RET64()*

# ***UNALIGNED\_PUT32()***

© 2004, QNX Software Systems Ltd.

*Write a misaligned 32-bit value safely*

## **Synopsis:**

```
#include <gulliver.h>

void UNALIGNED_PUT32(uint32_t *loc,
 uint32_t num);
```

## **Arguments:**

*loc*      The address where you want to write the value.

*num*      The value that you want to write.

## **Library:**

`libc`

## **Description:**

The *UNALIGNED\_PUT32()* macro lets you write the value *num* at the misaligned address *loc* without faulting.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*UNALIGNED\_PUT32()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16(), ENDIAN\_BE32(), ENDIAN\_BE64(),  
ENDIAN\_LE16(), ENDIAN\_LE32(), ENDIAN\_LE64(),  
ENDIAN\_RET32(), ENDIAN\_RET64(), ENDIAN\_SWAP16(),  
ENDIAN\_SWAP32(), ENDIAN\_SWAP64(), htonl(), htons(), ntohl(),  
ntohs(), UNALIGNED\_PUT16(), UNALIGNED\_PUT64(),  
UNALIGNED\_RET16(), UNALIGNED\_RET32(),  
UNALIGNED\_RET64()*

# ***UNALIGNED\_PUT64()***

© 2004, QNX Software Systems Ltd.

*Write a misaligned 64-bit value safely*

## **Synopsis:**

```
#include <gulliver.h>

void UNALIGNED_PUT64(uint64_t * loc,
 uint64_t num);
```

## **Arguments:**

*loc*      The address where you want to write the value.

*num*      The value that you want to write.

## **Library:**

`libc`

## **Description:**

The *UNALIGNED\_PUT64()* macro lets you write the value *num* at the misaligned address *loc* without faulting.



---

This macro isn't currently implemented.

---

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**Caveats:**

*UNALIGNED\_PUT64()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16(), ENDIAN\_BE32(), ENDIAN\_BE64(),  
ENDIAN\_LE16(), ENDIAN\_LE32(), ENDIAN\_LE64(),  
ENDIAN\_RET32(), ENDIAN\_RET64(), ENDIAN\_SWAP16(),  
ENDIAN\_SWAP32(), ENDIAN\_SWAP64(), htonl(), htons(), ntohl(),  
ntohs(), UNALIGNED\_PUT16(), UNALIGNED\_PUT32(),  
UNALIGNED\_RET16(), UNALIGNED\_RET32(),  
UNALIGNED\_RET64()*

# ***UNALIGNED\_RET16()***

© 2004, QNX Software Systems Ltd.

*Access a misaligned 16-bit value safely*

## **Synopsis:**

```
#include <gulliver.h>

uint16_t UNALIGNED_RET16(const uint16_t *loc);
```

## **Arguments:**

*loc*     The address where you want to get the value from.

## **Library:**

`libc`

## **Description:**

The *UNALIGNED\_RET16()* macro lets you access the misaligned 16-bit value pointed to by *loc* without faulting.

## **Returns:**

The 16-bit value pointed to by *loc*.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*UNALIGNED\_RET16()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16(), ENDIAN\_BE32(), ENDIAN\_BE64(),  
ENDIAN\_LE16(), ENDIAN\_LE32(), ENDIAN\_LE64(),  
ENDIAN\_RET32(), ENDIAN\_RET64(), ENDIAN\_SWAP16(),  
ENDIAN\_SWAP32(), ENDIAN\_SWAP64(), htonl(), htons(), ntohl(),  
ntohs(), UNALIGNED\_PUT16(), UNALIGNED\_PUT32(),  
UNALIGNED\_PUT64(), UNALIGNED\_RET32(),  
UNALIGNED\_RET64()*

# ***UNALIGNED\_RET32()***

© 2004, QNX Software Systems Ltd.

*Access a misaligned 32-bit value safely*

## **Synopsis:**

```
#include <gulliver.h>

uint32_t UNALIGNED_RET32(const uint32_t *loc);
```

## **Arguments:**

*loc*     The address where you want to get the value from.

## **Library:**

`libc`

## **Description:**

The *UNALIGNED\_RET32()* macro lets you access the misaligned 32-bit value pointed to by *loc* without faulting.

## **Returns:**

The 32-bit value pointed to by *loc*.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*UNALIGNED\_RET32()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16(), ENDIAN\_BE32(), ENDIAN\_BE64(),  
ENDIAN\_LE16(), ENDIAN\_LE32(), ENDIAN\_LE64(),  
ENDIAN\_RET32(), ENDIAN\_RET64(), ENDIAN\_SWAP16(),  
ENDIAN\_SWAP32(), ENDIAN\_SWAP64(), htonl(), htons(), ntohl(),  
ntohs(), UNALIGNED\_PUT16(), UNALIGNED\_PUT32(),  
UNALIGNED\_PUT64(), UNALIGNED\_RET16(),  
UNALIGNED\_RET64()*

# ***UNALIGNED\_RET64()***

© 2004, QNX Software Systems Ltd.

*Access a misaligned 64-bit value safely*

## **Synopsis:**

```
#include <gulliver.h>

uint64_t UNALIGNED_RET64(const uint64_t * loc);
```

## **Arguments:**

*loc*     The address where you want to get the value from.

## **Library:**

`libc`

## **Description:**

The *UNALIGNED\_RET64()* macro lets you access the misaligned 64-bit value pointed to by *loc* without faulting.



---

This macro isn't currently implemented.

---

## **Returns:**

The 64-bit value pointed to by *loc*.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*UNALIGNED\_RET64()* is implemented as a macro.

**See also:**

*ENDIAN\_BE16()*, *ENDIAN\_BE32()*, *ENDIAN\_BE64()*,  
*ENDIAN\_LE16()*, *ENDIAN\_LE32()*, *ENDIAN\_LE64()*,  
*ENDIAN\_RET32()*, *ENDIAN\_RET64()*, *ENDIAN\_SWAP16()*,  
*ENDIAN\_SWAP32()*, *ENDIAN\_SWAP64()*, *htonl()*, *htons()*, *ntohl()*,  
*ntohs()*, *UNALIGNED\_PUT16()*, *UNALIGNED\_PUT32()*,  
*UNALIGNED\_PUT64()*, *UNALIGNED\_RET16()*,  
*UNALIGNED\_RET32()*

# ***uname()***

© 2004, QNX Software Systems Ltd.

*Get information about the operating system*

## **Synopsis:**

```
#include <sys/utsname.h>

int uname(struct utsname * name);
```

## **Arguments:**

*name*     A pointer to a **utsname** where the function can store the information; see below.

## **Library:**

**libc**

## **Description:**

The *uname()* function stores information about the current operating system in the structure pointed to by the argument *name*.

The system name structure, **utsname**, is defined in **<sys/utsname.h>**, and contains at least the following structure members:

**char\*** *sysname*     The name of the OS.

**char\*** *nodename*     The name of this node.

**char\*** *release*     The current release level.

**char\*** *version*     The current version level.

**char\*** *machine*     The hardware type.

Each of these items is a null-terminated character array.



**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Examples:**

```

/*
 * The following program prints some information about the
 * system it's running on.
 */
#include <stdio.h>
#include <stdlib.h>
#include <sys/types.h>
#include <sys/utsname.h>

int main(void)
{
 struct utsname sysinfo;

 if(uname(&sysinfo) == -1) {
 perror("uname");
 return EXIT_FAILURE;
 }
 printf("system name : %s\n", sysinfo.sysname);
 printf("node name : %s\n", sysinfo.nodename);
 printf("release name : %s\n", sysinfo.release);
 printf("version name : %s\n", sysinfo.version);
 return EXIT_SUCCESS;
}

```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*

**uname** in the *Utilities Reference*

## Synopsis:

```
#include <stdio.h>

int ungetc(int c,
 FILE *fp);
```

## Arguments:

*c*     The character that you want to push back.

*fp*    The stream you want to push the character back on.

## Library:

libc

## Description:

The *ungetc()* function pushes the character specified by *c* back onto the input stream pointed to by *fp*. This character will be returned the next time that you read from the stream. The pushed-back character is discarded if you call *fflush()* or a file-positioning function (*fseek()*, *fsetpos()*, or *rewind()*) before performing the next read operation.

Only one character (the most recent one) of pushback is guaranteed.

The *ungetc()* function clears the end-of-file indicator, unless the value of *c* is EOF.

## Returns:

The character pushed back.

## Examples:

```
#include <stdio.h>
#include <stdlib.h>
#include <ctype.h>

int main(void)
{
 FILE *fp;
```

```
int c;
long value;

fp = fopen("file", "r");
value = 0;
c = fgetc(fp);
while(isdigit(c)) {
 value = value*10 + c - '0';
 c = fgetc(fp);
}
ungetc(c, fp); /* put last character back */
printf("Value=%ld\n", value);
fclose(fp);
return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*fopen()*, *getc()*, *getc\_unlocked()*, *ungetwc()*

## Synopsis:

```
#include <wchar.h>

wint_t ungetwc(wint_t wc,
 FILE * fp);
```

## Arguments:

*c*      The wide character that you want to push back.

*fp*     The stream you want to push the wide character back on.

## Library:

`libc`

## Description:

The *ungetwc()* function pushes the wide character specified by *wc* back onto the input stream pointed to by *fp*.

The pushed-back character will be returned the next time that you read from the stream but is discarded if you call *fflush()* or a file-positioning function (*fseek()*, *fsetpos()*, or *rewind()*) before the next read operation is performed.

Only one character (the most recent one) of pushback is guaranteed.

The *ungetwc()* function clears the end-of-file indicator, unless the value of *wc* is WEOF.

## Returns:

The character pushed back.

## Errors:

EILSEQ      Invalid character sequence or wide character.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*fopen()*, *getwc()*, *ungetc()*

## Synopsis:

```
#include <sys/socket.h>
#include <sys/un.h>

socket(AF_LOCAL,
 SOCK_STREAM,
 0);

socket(AF_LOCAL,
 SOCK_DGRAM,
 0);
```

## Description:

The UNIX-domain protocol family provides local (on-machine or QNX-network) interprocess communication through the normal *socket()* mechanisms. The UNIX-domain family supports the `SOCK_STREAM` and `SOCK_DGRAM` socket types and uses filesystem pathnames for addressing.

## Addressing

UNIX-domain addresses are variable-length filesystem pathnames of at most 104 characters. The `<sys/un.h>` include file defines this address:

```
struct sockaddr_un {
 u_char sun_len;
 u_char sun_family;
 char sun_path[104];
};
```

Binding a name to a UNIX-domain socket with *bind()* causes a socket file to be created in the filesystem. This file isn't removed when the socket is closed; you must use *unlink()* to remove the file.

You can use the macro *SUN\_LEN()* (defined in `<sys/un.h>`) to calculate the length of UNIX-domain address, required by *bind()* and *connect()*. The *sun\_path* field must be terminated by a NUL character

to be used with *SUN\_LEN()*, but the terminating NUL isn't part of the address.

The UNIX-domain protocol family doesn't support broadcast addressing or any form of "wildcard" matching on incoming messages. All addresses are absolute- or relative-pathnames of other UNIX-domain sockets. Normal filesystem access-control mechanisms are also applied when referencing pathnames (e.g. the destination of a *connect()* or *sendto()* must be writable).

## Protocols

The UNIX-domain protocol family consists of simple transport protocols that support the *SOCK\_STREAM* and *SOCK\_DGRAM* abstractions. UNIX-domain sockets also support the communication of QNX file descriptors through the use of the *msg\_control* field in the *msg* argument to *sendmsg()* and *recvmsg()*.

Any valid descriptor may be sent in a message. The file descriptor to be passed is described using a **struct cmsghdr** defined in the include file `<sys/socket.h>`. The type of the message is *SCM\_RIGHTS*, and the data portion of the messages is an array of integers representing the file descriptors to be passed. The number of descriptors being passed is defined by the length field of the message; the length field is the sum of the size of the header plus the size of the array of file descriptors.

The received descriptor is a duplicate of the sender's descriptor, as if it were created with a call to *dup()*. Descriptors awaiting delivery or purposely not received are automatically closed by the system when the destination socket is closed.

## LOCAL\_CREDS

There is one socket-level option for *setsockopt()* and *getsockopt()* available in the UNIX-domain. The *LOCAL\_CREDS* option may be enabled on a *SOCK\_DGRAM* or a *SOCK\_STREAM* socket. This option provides a mechanism for the receiver to receive the credentials of the process as a *recvmsg()* message. The *msg\_control* field in the **msg\_hdr** structure points to a buffer that contains a **cmsghdr** structure



followed by a variable length `sockcred` structure defined in `<sys/socket.h>` as follows:

```
struct sockcred {
 uid_t sc_uid; /* real user id */
 uid_t sc_euid; /* effective user id */
 gid_t sc_gid; /* real group id */
 gid_t sc_egid; /* effective group id */
 int sc_ngroups; /* number of supplemental groups */
 gid_t sc_groups[1]; /* variable length */
};
```

The `SOCKCREDSIZE()` macro computes the size of the `sockcred` structure for a specified number of groups. The `cmsghdr` fields have the following values:

```
msg_len = sizeof(struct cmsghdr) + SOCKCREDSIZE(ngroups)
msg_level = SOL_SOCKET
msg_type = SCM_CREDS
```

### See also:

*bind()*, *connect()*, *dup()*, *getsockopt()*, *recvmsg()*, *sendmsg()*, *sendto()*, *setsockopt()*, *socket()*, *unlink()*

# ***unlink()***

© 2004, QNX Software Systems Ltd.

*Remove a link to a file*

## **Synopsis:**

```
#include <unistd.h>

int unlink(const char * path);
```

## **Arguments:**

*path*     The name of the file that you want to unlink.

## **Library:**

`libc`

## **Description:**

The *unlink()* function removes a link to a file:

- If the *path* names a symbolic link, *unlink()* removes the link, but doesn't affect the file or directory that the link goes to.
- If the *path* isn't a symbolic link, *unlink()* removes the link and decrements the link count of the file that the link refers to.

If the link count of the file becomes zero, and no process has the file open, then the space that the file occupies is freed, and no one can access the file anymore.

If one or more processes have the file open when the last link is removed, the link is removed, but the removal of the file is delayed until all references to it have been closed.

This function is equivalent to *remove()*.



---

To remove a directory, call *rmdir()*.

---

**Returns:**

- |         |                                              |
|---------|----------------------------------------------|
| 0       | The operation was successful.                |
| Nonzero | The operation failed ( <i>errno</i> is set). |

**Errors:**

- |              |                                                                                                                                                                                                                    |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EACCES       | Search permission is denied for a component of <i>path</i> , or write permission is denied on the directory containing the link to be removed.                                                                     |
| EBUSY        | The directory named by <i>path</i> cannot be unlinked because it's being used by the system or another process, and the target filesystem or resource manager considers this to be an error.                       |
| ENAMETOOLONG | The <i>path</i> argument exceeds PATH_MAX in length, or a pathname component is longer than NAME_MAX.                                                                                                              |
| ENOENT       | The named file doesn't exist, or <i>path</i> is an empty string.                                                                                                                                                   |
| ENOSYS       | The <i>unlink()</i> function isn't implemented for the filesystem specified by <i>path</i> .                                                                                                                       |
| ENOTDIR      | A component of <i>path</i> isn't a directory.                                                                                                                                                                      |
| EPERM        | The file named by <i>path</i> is a directory, and either the calling process doesn't have the appropriate privileges, or the target filesystem or resource manager prohibits using <i>unlink()</i> on directories. |
| EROFS        | The directory entry to be unlinked resides on a read-only filesystem.                                                                                                                                              |

**Examples:**

```
#include <unistd.h>
#include <stdlib.h>

int main(void)
{
 if(unlink("vm.tmp")) {
 puts("Error removing vm.tmp!");
 return EXIT_FAILURE;
 }

 return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*chdir(), chmod(), close(), errno, getcwd(), link(), mkdir(), open(), pathmgr\_symlink(), pathmgr\_unlink(), remove(), rename(), rmdir(), stat(), symlink()*

**Synopsis:**

```
#include <stdlib.h>

void unsetenv(const char* name);
```

**Arguments:**

*name*      The name of the environment variable that you want to delete.

**Library:**

libc

**Description:**

The *unsetenv()* function removes the environment variable named *name* from the process's environment.

**Classification:**

POSIX 1003.1a

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

The *unsetenv()* function manipulates the environment pointed to by the global *environ* variable.

**See also:**

*clearenv(), getenv(), putenv(), setenv()*

## Synopsis:

```
#include <unistd.h>

int usleep(useconds_t useconds);
```

## Arguments:

*useconds*      The number of microseconds that you want to process to sleep for. This must be less than 1,000,000.

## Library:

libc

## Description:

The *usleep()* function suspends the calling thread until *useconds* microseconds of realtime have elapsed, or until a signal that isn't ignored is received. The time spent suspended could be longer than the requested amount due to the scheduling of other, higher-priority threads.

If *useconds* is 0, *usleep()* has no effect.

## Returns:

0      Success.

-1      An error occurred (*errno* is set).

## Errors:

EAGAIN      No timer resources are available to satisfy the request.

EINVAL      The *useconds* argument is too large.

## Examples:

```
/*
 * The following program sleeps for the
 * number of microseconds specified in argv[1].
 */
#include <stdlib.h>
#include <unistd.h>

int main(int argc, char **argv)
{
 useconds_t microseconds;

 microseconds = (useconds_t)strtol(argv[1], NULL, 0);
 if(usleep(microseconds) == 0) {
 return EXIT_SUCCESS;
 }
 return EXIT_FAILURE;
}
```

## Classification:

Standard Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*alarm()*, *nanosleep()*, *sigaction()*, *sleep()*, *timer\_create()*,  
*timer\_delete()*, *timer\_getoverrun()*, *timer\_gettime()*, *timer\_settime()*,  
*ualarm()*



**Synopsis:**

```
#include <sys/types.h>
#include <utime.h>

struct utimbuf {
 time_t actime; /* access time */
 time_t modtime; /* modification time */
};

int utime(const char* path,
 const struct utimbuf* times);
```

**Arguments:**

*path*      The path name for the file whose modification time you want to get or set.

*times*     NULL, or a pointer to a `utimbuf` structure where the function can store the modification time.

**Library:**

`libc`

**Description:**

The `utime()` function records the modification time for the file or directory identified by *path*.

If the *times* argument is NULL, the access and modification times of the file or directory are set to the current time. The effective user ID of the process must match the owner of the file or directory, or the process must have write permission to the file or directory, or appropriate privileges in order to use the `utime()` function in this way.

If the *times* argument isn't NULL, its interpreted as a pointer to a `utimbuf` structure, and the access and modification times of the file or directory are set to the values contained in the designated structure. Only the owner of the file or directory, and processes with appropriate

privileges are permitted to use the *utime()* function in this way. The access and modification times are taken from the *actime* and *modtime* fields in this structure.

**Returns:**

- 0 Success.
- 1 An error occurred; *errno* is set.

**Errors:**

- EACCES Search permission is denied for a component of *path*, or the *times* argument is NULL, and the effective user ID of the process doesn't match the owner of the file, and write access is denied.
- ENAMETOOLONG  
The argument *path* exceeds PATH\_MAX in length, or a pathname component is longer than NAME\_MAX.
- ENOENT The specified *path* doesn't exist, or *path* is an empty string.
- ENOTDIR A component of *path* isn't a directory.
- EPERM The *times* argument isn't NULL, and the calling process's effective user ID has write access to the file but doesn't match the owner of the file, and the calling process doesn't have the appropriate privileges.
- EROFS The named file resides on a read-only filesystem.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <utime.h>

int main(int argc, char *argv[])
{
 if((utime(argv[1], NULL) != 0) && (argc > 1)) {
```

```
 printf("Unable to set time for %s\n", argv[1]);
 }
 return EXIT_SUCCESS;
}
```

## Classification:

POSIX 1003.1

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*, *futime()*

# ***utimes()***

© 2004, QNX Software Systems Ltd.

*Set a file's access and modification times*

## **Synopsis:**

```
#include <sys/time.h>

int utimes(const char * __path,
 const struct timeval * __times);
```

## **Arguments:**

- \_\_path*      The name of the files whose times you want to set.
- \_\_times*     NULL, or an array of **timeval** structures:
- The first array member represents the date and time of last access.
  - The second member represents the date and time of last modification.

## **Library:**

**libc**

## **Description:**

The *utimes()* function sets the access and modification times of the file pointed to by the *\_\_path* argument to the value of the *\_\_times* argument. This function allows time specifications accurate to the microsecond.

The times in the **timeval** structure are measured in seconds and microseconds since the Unix Epoch (00:00:00 January 1, 1970 Coordinated Universal Time (UTC)), although rounding toward the nearest second may occur.

If the *\_\_times* argument is NULL, the access and modification times of the file are set to the current time. The effective user ID of the process must be the same as the owner of the file, or must have write access to the file or superuser privileges to use this call in this manner. On completion, *utimes()* marks the time of the last file status change, *st\_ctime*, for update.

**Returns:**

- 0 Success.
- 1 An error occurred (*errno* is set).

**Errors:**

- EACCES Search permission is denied by a component of the path prefix; or the *\_times* argument is NULL and the effective user ID of the process doesn't match the owner of the file and write access is denied.
- EFAULT The *\_path* or *\_times* argument points to an illegal address.
- EINTR A signal was caught during the *utimes()* function.
- EINVAL The number of microseconds specified in one or both of the *timeval* structures pointed to by *\_times* was greater than or equal to 1,000,000 or less than 0.
- EIO An I/O error occurred while reading from or writing to the filesystem.
- ELOOP Too many symbolic links were encountered in resolving *\_path*.
- EMULTIHOP Components of *\_path* require hopping to multiple remote machines and the filesystem doesn't allow it.
- ENAMETOOLONG  
The length of the *\_path* argument exceeds PATH\_MAX or a pathname component is longer than NAME\_MAX.
- ENOLINK The *\_path* argument points to a remote machine and the link to that machine is no longer active.

|              |                                                                                                                                                                                                                          |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ENOENT       | A component of <i>_path</i> doesn't name an existing file or <i>_path</i> is an empty string.                                                                                                                            |
| ENOTDIR      | A component of the path prefix isn't a directory.                                                                                                                                                                        |
| EPERM        | The <i>_times</i> argument isn't NULL and the calling process's effective user ID has write access to the file but doesn't match the owner of the file, and the calling process doesn't have the appropriate privileges. |
| EROFS        | The filesystem containing the file is read-only.                                                                                                                                                                         |
| ENAMETOOLONG | Path name resolution of a symbolic link produced an intermediate result whose length exceeds PATH_MAX.                                                                                                                   |

**Classification:**

Legacy Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*stat()*

**Synopsis:**

```
struct utmp {
 char ut_user[UT_NAMESIZE];
#define ut_name ut_user
 char ut_id[4];
 char ut_line[UT_LINESIZE];
 pid_t ut_pid;
 short ut_type;
 struct exit_status {
 short e_termination;
 short e_exit;
 } ut_exit;
 short ut_spare;
 time_t ut_time;
};
```

**Description:**

The **utmp** structure describes an entry in a user-information file. The members include:

- |                |                                                                                                                                                                                                                                |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>ut_user</i> | The user's login name.                                                                                                                                                                                                         |
| <i>ut_id</i>   | The line number.                                                                                                                                                                                                               |
| <i>ut_line</i> | The device name (console).                                                                                                                                                                                                     |
| <i>ut_pid</i>  | The process ID.                                                                                                                                                                                                                |
| <i>ut_type</i> | The type of entry. The possible values are: <ul style="list-style-type: none"><li>• EMPTY</li><li>• RUN_LVL</li><li>• BOOT_TIME</li><li>• OLD_TIME</li><li>• NEW_TIME</li><li>• INIT_PROCESS</li><li>• LOGIN_PROCESS</li></ul> |

- USER\_PROCESS
- DEAD\_PROCESS
- ACCOUNTING

*ut\_exit* The exit status of a process marked as DEAD\_PROCESS. The structure `exit_status` includes at least the following members:

- *e\_termination* — the termination status.
- *e\_exit* — the exit status.

*ut\_time* The time that this entry was made.

## Classification:

Unix

## See also:

*endutent()*, *getutent()*, *getutid()*, *getutline()*, *pututline()*, *setutent()*, *utmpname()*

**login** in the *Utilities Reference*



**Synopsis:**

```
#include <utmp.h>

void utmpname(char * __filename);
```

**Arguments:**

*\_\_filename*     The new filename that you want to use.

**Library:**

libc

**Description:**

The *utmpname()* function lets you change the name of the file examined from the default file (*\_PATH\_UTMP*) to any other file. If the file doesn't exist, this won't be apparent until the first attempt to reference the file is made. This function doesn't open the file. It just closes the old file if it's currently open and saves the new file name.

**Files:**

*\_PATH\_UTMP*  
Specifies the user information file.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*endutent()*, *getutent()*, *getutid()*, *getutline()*, *pututline()*, *setutent()*,  
**utmp**

**login** in the *Utilities Reference*

## Synopsis:

```
#include <stdlib.h>

char* utoa(unsigned int value,
 char* buffer,
 int radix);
```

## Arguments:

*value*      The value to convert into a string.

*buffer*     A buffer in which the function stores the string. The size of the buffer must be at least:

$8 \times \text{sizeof}( \text{int} ) + 1$

bytes when converting values in base 2 (binary).

*radix*      The base to use when converting the number.

## Library:

`libc`

## Description:

The *utoa()* function converts the unsigned binary integer *value* into the equivalent string in base *radix* notation, storing the result in the character array pointed to by *buffer*. A null character is appended to the result.

## Returns:

A pointer to the result.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

int main(void)
{
 int base;
 char buffer[18];

 for(base = 2; base <= 16; base = base + 2)
 printf("%2d %s\n", base,
 utoa((unsigned) 12765, buffer, base));
 return EXIT_SUCCESS;
}
```

produces the output:

```
2 11000111011101
4 3013131
6 135033
8 30735
10 12765
12 7479
14 491b
16 31dd
```

**Classification:**

QNX 4

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*atoi(), atol(), itoa(), ltoa(), sscanf(), strtol(), strtoul(), ultoa()*

## ***va\_arg()***

© 2004, QNX Software Systems Ltd.

*Get the next item in a list of variable arguments*

### **Synopsis:**

```
#include <stdarg.h>

type va_arg(va_list param,
 type);
```

### **Arguments:**

*param*     The **va\_list** object that you initialized with the *va\_start()* macro.

*type*       The type of the next argument.

### **Library:**

**libc**

### **Description:**

You can use the *va\_arg()* macro to get the next argument in a list of variable arguments.



---

**CAUTION:** Take special care when using varargs on some platforms; see “Varargs and coercion,” below.

---

You must use *va\_arg()* with the associated macros *va\_copy()*, *va\_start()* and *va\_end()*. A sequence such as:

```
void example(char *dst, ...)
{
 va_list curr_arg;
 int next_arg;

 va_start(curr_arg, dst);
 next_arg = va_arg(curr_arg, int);
 :
}
```

causes *next\_arg* to be assigned the value of the next variable argument. The argument *type* (which is `int` in the example) is the type of the argument originally passed to the function.



The last argument before the ellipsis (`...`) has to be an `int` or a type that doesn't change in size if cast to an `int`. If the argument is promoted, the ANSI/ISO standard says the behavior is undefined, and so depends on the compiler and the library.

You must execute the macro `va_start()` first, in order to initialize the variable *curr\_arg* properly, and execute the macro `va_end()` after getting all the arguments.

The data item *curr\_arg* is of type `va_list` that contains the information to permit successive acquisitions of the arguments.

The following functions use a “varargs” list:

|                          |                          |                         |
|--------------------------|--------------------------|-------------------------|
| <code>verr()</code>      | <code>vscanf()</code>    | <code>vsyslog()</code>  |
| <code>verrx()</code>     | <code>vslogf()</code>    | <code>vwarn()</code>    |
| <code>vfprintf()</code>  | <code>vsnprintf()</code> | <code>vwarnx()</code>   |
| <code>vfscanf()</code>   | <code>vsprintf()</code>  | <code>vwprintf()</code> |
| <code>vfwprintf()</code> | <code>vsscanf()</code>   | <code>vwscanf()</code>  |
| <code>vwscanf()</code>   | <code>vswprintf()</code> |                         |
| <code>vprintf()</code>   | <code>vswscanf()</code>  |                         |

## Varargs and coercion

On some platforms, such as PowerPC, the `va_list` type is an array; on other platforms, such as x86, it isn't. This can lead to problems.

Consider the following example. It seems correct, but on PowerPC platforms, it doesn't print 2:

```
#include <stdio.h>
#include <stdarg.h>

void handle_foo(char *fmt, va_list *pva) {
 printf("%d\n", va_arg(*pva, int));
}

void vfoo(char *fmt, va_list va) {
```

```
 handle_foo(fmt, &va);
}

void foo(char *fmt, ...) {
 va_list va;

 va_start(va, fmt);
 vfoo(fmt, va);
 va_end(va);
}

int main() {
 foo("", 2);
 return 0;
}
```

The C standard says that prototypes such as *vfoo()* have the array type silently coerced to be a pointer to a base type. This makes things work when you pass an array object to the function. An array-typed expression is converted to a pointer to the first element when used in an rvalue context, so the coercion in the function makes everybody happy.

The problem occurs when you then pass the address of the `va_list` parameter to another function. The function expects a pointer to the array, but what it *really* gets is a pointer to a pointer (because of the original conversion). If you use the `va_list` type in the second function, you won't get the right data.

Here's the example modified so that it works in all cases:

```
#include <stdio.h>
#include <stdarg.h>

void handle_foo(char *fmt, va_list *pva) {
 printf("%d\n", va_arg(*pva, int));
}

void vfoo(char *fmt, va_list va) {
 va_list temp;

 va_copy(temp, va);
 handle_foo(fmt, &temp);
 va_end(temp);
}
```



```

void foo(char *fmt, ...) {
 va_list va;

 va_start(va, fmt);
 vfoo(fmt, va);
 va_end(va);
}

int main() {
 foo("", 2);
 return 0;
}

```

Using *va\_copy()* “undoes” the coercion that happens in the parameter list, so that *handle\_foo()* gets the proper data.

## Returns:

The value of the next variable argument, according to type passed as the second parameter.

## Examples:

```

#include <stdio.h>
#include <stdarg.h>
#include <stdlib.h>

static void test_fn(const char *msg,
 const char *types,
 ...);

int main(void)
{
 printf("VA...TEST\n");
 test_fn("PARAMETERS: 1, \"abc\", 546",
 "isi", 1, "abc", 546);
 test_fn("PARAMETERS: \"def\", 789",
 "si", "def", 789);
 return EXIT_SUCCESS;
}

static void test_fn(
 const char *msg, /* message to be printed */
 const char *types, /* parameter types (i,s) */
 ...) /* variable arguments */
{
 va_list argument;
 int arg_int;

```

```

char *arg_string;
const char *types_ptr;

types_ptr = types;
printf("\n%s -- %s\n", msg, types);
va_start(argument, types);
while(*types_ptr != '\0') {
 if (*types_ptr == 'i') {
 arg_int = va_arg(argument, int);
 printf("integer: %d\n", arg_int);
 } else if (*types_ptr == 's') {
 arg_string = va_arg(argument, char *);
 printf("string: %s\n", arg_string);
 }
 ++types_ptr;
}
va_end(argument);
}

```

produces the output:

```

VA...TEST

PARAMETERS: 1, "abc", 546 -- isi
integer: 1
string: abc
integer: 546

PARAMETERS: "def", 789 -- si
string: def
integer: 789

```

## Classification:

ANSI

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*va\_arg()* is a macro.

**See also:**

*va\_copy()*, *va\_end()*, *va\_start()*

## ***va\_copy()***

© 2004, QNX Software Systems Ltd.

*Make a copy of a variable argument list*

### **Synopsis:**

```
#include <stdarg.h>

void va_copy(va_list d,
 va_list s);
```

### **Arguments:**

- d* A **va\_list** object into which you want to copy the list.
- s* The **va\_list** object that you initialized with the *va\_start()* macro and that you want to copy.

### **Library:**

**libc**

### **Description:**

The *va\_copy()* macro creates a copy of a list of variable arguments.

You can use the *va\_copy()* macro with the associated macros *va\_arg()*, *va\_start()*, and *va\_end()*, especially to avoid problems on some platforms. For more information, see “Varargs and coercion” in the documentation for *va\_arg()*.

### **Examples:**

See *va\_arg()*.

### **Classification:**

ANSI

#### **Safety**

---

Cancellation point No

Interrupt handler Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**Caveats:**

*va\_copy()* is a macro.

**See also:**

*va\_arg()*, *va\_end()*, *va\_start()*

## ***va\_end()***

© 2004, QNX Software Systems Ltd.

*Finish getting items from a variable argument list*

### **Synopsis:**

```
#include <stdarg.h>

void va_end(va_list param);
```

### **Arguments:**

*param*     The **va\_list** object that you initialized with the *va\_start()* macro.

### **Library:**

libc

### **Description:**

Use the *va\_end()* macro to complete the acquisition of arguments from a list of variable arguments. You must use it with the associated macros *va\_copy()*, *va\_start()*, and *va\_arg()*. For more information, see *va\_arg()*.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*va\_end()* is a macro.

**See also:**

*va\_arg()*, *va\_copy()*, *va\_start()*

## ***va\_start()***

© 2004, QNX Software Systems Ltd.

*Start getting items from a variable argument list*

---

### **Synopsis:**

```
#include <stdarg.h>

void va_start(va_list param,
 previous);
```

### **Arguments:**

*param*      A `va_list` object that the “varargs” macros can use to locate the arguments.

*previous*    The argument that immediately precedes the “...” notation in the original function definition.

### **Library:**

`libc`

### **Description:**

Use the `va_start()` macro to start the acquisition of arguments from a list of variable arguments.

You must use the `va_start()` macro with the associated macros `va_arg()`, `va_copy()`, and `va_end()`. For each call to `va_start()`, you must have a matching call to `va_end()`. For more information, see `va_arg()`.

### **Examples:**

See `va_arg()`.

### **Classification:**

ANSI



**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**Caveats:**

*va\_start()* is a macro.

**See also:**

*va\_arg()*, *va\_copy()*, *va\_end()*

# ***valloc()***

© 2004, QNX Software Systems Ltd.

*Allocate a heap block aligned on a page boundary*

## **Synopsis:**

```
#include <stdarg.h>

void * valloc(size_t size);
```

## **Arguments:**

*size*     The size of the block to allocate, in bytes.

## **Library:**

`libc`

## **Description:**

The *valloc()* function allocates a heap block that's aligned on a page boundary. It's equivalent to:

```
memalign(sysconf(_SC_PAGESIZE), size);
```

## **Returns:**

See *memalign()*.

## **Classification:**

QNX Neutrino

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memalign()*, *sysconf()*

## ***verr()*, *verrx()***

© 2004, QNX Software Systems Ltd.

*Display a formatted error message, and then exit (varargs)*

### **Synopsis:**

```
#include <err.h>

void verr(int eval,
 const char *fmt,
 va_list args);

void verrx(int eval,
 const char *fmt,
 va_list args);
```

### **Arguments:**

*eval*     The value to use as the exit code of the process.

*fmt*     NULL, or a *printf()*-style string used to format the message.

*args*     A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *err()* and *warn()* family of functions display a formatted error message on *stderr*. For a comparison of the members of this family, see *err()*.

The *verr()* function produces a message that consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, followed by a colon and a space, if the *fmt* argument isn't NULL
- the string associated with the current value of *errno*
- a newline character.

The *verrx()* function produces a similar message, except that it doesn't include the string associated with *errno*. The message consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, if the *fmt* argument isn't NULL
- a newline character.

The *verr()* and *verrx()* functions don't return, but exit with the value of the argument *eval*.

## Classification:

Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*err()*, *errx()*, *stderr*, *strerror()*, *vwarn()*, *vwarnx()*, *warn()*, *warnx()*

# ***vfork()***

© 2004, QNX Software Systems Ltd.

*Spawn a new process and block the parent*

## **Synopsis:**

```
#include <process.h>

pid_t vfork(void);
```

## **Library:**

libc

## **Description:**

This function spawns a new process and blocks the parent until the child process calls *execve()* or exits (by calling *\_exit()* or abnormally).

## **Returns:**

A value of zero to the child process, and (later) the child's process ID in the parent. If an error occurs, no child process is created, and the function returns -1 and sets *errno*.

## **Errors:**

|        |                                                                                                                                                     |
|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The system-imposed limit on the total number of processes under execution would be exceeded. This limit is determined when the system is generated. |
| ENOMEM | There isn't enough memory for the new process.                                                                                                      |

## **Classification:**

Standard Unix

### **Safety**

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |    |
|----------------|----|
| Signal handler | No |
| Thread         | No |

**Caveats:**

To avoid a possible deadlock situation, processes that are children in the middle of a *vfork()* are never sent SIGTTOU or SIGTTIN signals; rather, output or ioctls are allowed and input attempts result in an EOF indication.

**See also:**

*execve()*, *\_exit()*, *fork()*, *ioctl()*, *sigaction()*, *wait()*

## ***vfprintf()***

© 2004, QNX Software Systems Ltd.

*Write formatted output to a file (varargs)*

### **Synopsis:**

```
#include <stdio.h>
#include <stdarg.h>

int vfprintf(FILE* fp,
 const char* format,
 va_list arg);
```

### **Arguments:**

|               |                                                                                                                                                                               |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>fp</i>     | The stream to which you want to send the output.                                                                                                                              |
| <i>format</i> | A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see <i>printf()</i> . |
| <i>arg</i>    | A variable-argument list of the additional arguments, which you must have initialized with the <i>va_start()</i> macro.                                                       |

### **Library:**

`libc`

### **Description:**

The *vfprintf()* function writes output to the file pointed to by *fp*, under control of the argument *format*.

The *vfprintf()* function is a “varargs” version of *fprintf()*.

### **Returns:**

The number of characters written, or a negative value if an output error occurred (*errno* is set).

### **Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>
```



```
FILE *LogFile;

/* a general error routine */

void errmsg(char *format, ...)
{
 va_list arglist;

 fprintf(stderr, "Error: ");
 va_start(arglist, format);
 vfprintf(stderr, format, arglist);
 va_end(arglist);
 if(LogFile != NULL) {
 fprintf(LogFile, "Error: ");
 va_start(arglist, format);
 vfprintf(LogFile, format, arglist);
 va_end(arglist);
 }
}

int main(void)
{
 LogFile = fopen("error.log", "w");
 errmsg("%s %d %s", "Failed", 100, "times");
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

### **See also:**

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*,  
*va\_start()*, *vfwprintf()*, *vprintf()*, *vsprintf()*, *vswprintf()*,  
*vwprintf()*, *wprintf()*

## Synopsis:

```
#include <stdio.h>
#include <stdarg.h>

int vfscanf(FILE *fp,
 const char *format,
 va_list arg);
```

## Arguments:

- fp*            The stream that you want to read from.
- format*        A string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.
- arg*            A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

## Library:

`libc`

## Description:

The *vfscanf()* function scans input from the file designated by *fp*, under control of the argument *format*.

The *vfscanf()* function is a “varargs” version of *fscanf()*.

## Returns:

The number of input arguments for which values were successfully scanned and stored, or EOF when the scanning is stopped by reaching the end of the input stream before storing any values.

**Errors:**

If an error occurs, *errno* indicates the type of error.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>

void ffind(FILE *fp, char *format, ...)
{
 va_list arglist;

 va_start(arglist, format);
 vfscanf(fp, format, arglist);
 va_end(arglist);
}

int main(void)
{
 int day, year;
 char weekday[10], month[12];

 ffind(stdin,
 "%s %s %d %d",
 weekday, month, &day, &year);
 printf("\n%s, %s %d, %d\n",
 weekday, month, day, year);
 return EXIT_SUCCESS;
}
```

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno*, *fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*,  
*vwscanf()*, *vscanf()*, *vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

## ***vfwprintf()***

© 2004, QNX Software Systems Ltd.

*Write formatted wide-character output to a file (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vfwprintf(FILE * fp,
 const wchar_t * format,
 va_list arg);
```

### **Arguments:**

*fp*            The stream to which you want to send the output.

*format*        A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

*arg*           A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vfwprintf()* function writes output to the file pointed to by *fp*, under control of the argument *format*.

The *vfwprint()* function is the wide-character version of *vfprintf()*, and is a “varargs” version of *fwprintf()*.

### **Returns:**

The number of wide characters written, excluding the terminating **NUL**, or a negative number if an error occurred (*errno* is set).

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*,  
*va\_start()*, *vfprintf()*, *vprintf()*, *vsprintf()*, *vsnprintf()*, *vswprintf()*,  
*vwprintf()*, *wprintf()*

## ***vfwscanf()***

© 2004, QNX Software Systems Ltd.

*Scan input from a file (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vfwscanf(FILE * fp,
 const wchar_t *format,
 va_list arg);
```

### **Arguments:**

- fp*            The stream that you want to read from.
- format*        A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.
- arg*            A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vfwscanf()* function scans input from the file designated by *fp*, under control of the argument *format*.

The *vfwscanf()* function is the wide-character version of *vscanf()*, and is a “varargs” version of *fwscanf()*.

### **Returns:**

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning reached the end of the input stream before storing any values.



## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*,  
*vfprintf()*, *vscanf()*, *vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

## ***vprintf()***

© 2004, QNX Software Systems Ltd.

*Write formatted output to standard output (varargs)*

### **Synopsis:**

```
#include <stdio.h>
#include <stdarg.h>

int vprintf(const char* format,
 va_list arg);
```

### **Arguments:**

*format*     A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

*arg*        A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vprintf()* function writes output to the file *stdout*, under control of the argument *format*.

The *vprintf()* function is a “varargs” version of *printf()*.

### **Returns:**

The number of characters written, or a negative value if an output error occurred (*errno* is set).

### **Examples:**

Use *vprintf()* in a general error message routine:

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>
```

```
void errmsg(char* format, ...)
{
 va_list arglist;

 printf("Error: ");
 va_start(arglist, format);
 vprintf(format, arglist);
 va_end(arglist);
}

int main(void)
{
 errmsg("%s %d %s", "Failed", 100, "times");
 return EXIT_SUCCESS;
}
```

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## See also:

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*,  
*va\_start()* *vfprintf()*, *vwprintf()*, *vsnprintf()*, *vsprintf()*, *vswprintf()*,  
*vwprintf()*, *wprintf()* *vsprintf()*

## ***vscanf()***

© 2004, QNX Software Systems Ltd.

*Scan input from standard input (varargs)*

### **Synopsis:**

```
#include <stdio.h>
#include <stdarg.h>

int vscanf(const char * format,
 va_list args);
```

### **Arguments:**

*format*     A string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

*args*       A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vscanf()* function scans input from *stdin*, under control of the argument *format*. For information about the *format* string, see the description of *scanf()*.

The *vscanf()* function is a “varargs” version of *scanf()*.

### **Returns:**

The number of input arguments for which values were successfully scanned and stored, or EOF when the scanning is stopped by reaching the end of the input stream before storing any values.

### **Errors:**

If an error occurs, *errno* indicates the type of error.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>

void find(char *format, ...)
{
 va_list arglist;

 va_start(arglist, format);
 vscanf(format, arglist);
 va_end(arglist);
}

int main(void)
{
 int day, year;
 char weekday[10], month[12];

 ffind("%s %s %d %d",
 weekday, month, &day, &year);
 printf("\n%s, %s %d, %d\n",
 weekday, month, day, year);
 return EXIT_SUCCESS;
}
```

**Classification:**

Unix

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno*, *fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*,  
*vfscanf()*, *vwscanf()*, *vsscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

**Synopsis:**

```
#include <stdio.h>
#include <sys/slog.h>

int vslogf(int opcode,
 int severity,
 const char * fmt,
 va_list arg);
```

**Arguments:**

- opcode*      A combination of a *major* and *minor* code. Create the *opcode* using the `_SLOG_SETCODE(major, minor)` macro that's defined in `<sys/slog.h>`.  
The *major* and *minor* codes are defined in `<sys/slogcodes.h>`.
- severity*    The severity of the log message; see "Severity levels," in the documentation for `slogf()`.
- fmt*         A standard `printf()` string.
- arg*         A variable-argument list of the additional arguments, which you must have initialized with the `va_start()` macro.

**Library:**

`libc`

**Description:**

The `slog*()` functions send log messages to the system logger, `slogger`. To send formatted messages, use `vslogf()`. If you have programs that scan log files for specified codes, you can use `slogb()` or `slogi()` to send a block of structures or `int`'s, respectively.

This function is a "varargs" version of `slogf()`.

**Classification:**

QNX Neutrino

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*slogb()*, *slogi()*, *slogf()*

**slogger**, **sloginfo** in the *Utilities Reference*



*Write formatted output to a character array, up to a given maximum number of characters (varargs)*

### **Synopsis:**

```
#include <stdarg.h>
#include <stdio.h>

int vsnprintf(char* buf,
 size_t count,
 const char* format,
 va_list arg);
```

### **Arguments:**

- buf*            A pointer to the buffer where you want to function to store the formatted string.
- count*         The maximum number of characters to store in the buffer, including a terminating null character.
- format*        A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.
- arg*            A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vsnprintf()* function formats data under control of the *format* control string and stores the result in *buf*. The maximum number of characters to store, including a terminating null character, is specified by *count*.

The *vsnprintf()* function is a “varargs” version of *sprintf()*.

**Returns:**

The number of characters that would have been written into the array, not counting the terminating null character, had *count* been large enough. It does this even if *count* is zero; in this case *buf* can be NULL.

If an error occurred, *vsnprintf()* returns a negative value and sets *errno*.

**Examples:**

Use *vsnprintf()* in a general error message routine:

```
#include <stdio.h>
#include <stdarg.h>
#include <string.h>

char msgbuf[80];

char *fmtmsg(char *format, ...)
{
 va_list arglist;

 va_start(arglist, format);
 strcpy(msgbuf, "Error: ");
 vsnprintf(&msgbuf[7], 80-7, format, arglist);
 va_end(arglist);
 return(msgbuf);
}

int main(void)
{
 char *msg;

 msg = fmtmsg("%s %d %s", "Failed", 100, "times");
 printf("%s\n", msg);

 return 0;
}
```

## Classification:

Standard Unix

### **Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Yes                     |

## Caveats:

It's safe to call *vsnprintf()* in a signal handler if the data isn't floating point.

## See also:

*errno*, *fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*, *va\_start()*, *vfprintf()*, *vwprintf()*, *vprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*, *wprintf()*

## ***vsprintf()***

© 2004, QNX Software Systems Ltd.

*Write formatted output to a buffer (varargs)*

### **Synopsis:**

```
#include <stdio.h>
#include <stdarg.h>

int vsprintf(char* buf,
 const char* format,
 va_list arg);
```

### **Arguments:**

*buf*            A pointer to the buffer where you want the function to store the formatted string.

*format*        A string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

*arg*            A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vsprintf()* function formats data under control of the *format* control string, and writes the result to *buf*.

The *vsprintf()* function is a “varargs” version of *sprintf()*.

### **Returns:**

The number of characters written, or a negative value if an output error occurred (*errno* is set).

**Examples:**

Use *vsprintf()* in a general error message routine:

```
#include <stdio.h>
#include <stdarg.h>
#include <stdlib.h>
#include <string.h>

char msgbuf[80];

char *fmtmsg(char *format, ...)
{
 va_list arglist;

 va_start(arglist, format);
 strcpy(msgbuf, "Error: ");
 vsprintf(&msgbuf[7], format, arglist);
 va_end(arglist);
 return(msgbuf);
}

int main(void)
{
 char *msg;

 msg = fmtmsg("%s %d %s", "Failed", 100, "times");
 printf("%s\n", msg);
 return EXIT_SUCCESS;
}
```

**Classification:**

ANSI

**Safety**


---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Yes                     |

### **Caveats:**

It's safe to call *vsprintf()* in a signal handler if the data isn't floating point.

### **See also:**

*fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*, *va\_start()*,  
*vfprintf()*, *vwprintf()*, *vprintf()*, *vsprintf()*, *vswprintf()*, *vwprintf()*,  
*wprintf()*

## Synopsis:

```
#include <stdio.h>
#include <stdarg.h>

int vsscanf(const char* in_string,
 const char* format,
 va_list arg);
```

## Arguments:

|                  |                                                                                                                                                                             |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>in_string</i> | The string that you want to read from.                                                                                                                                      |
| <i>format</i>    | A string that specifies the format of the input. For more information, see <i>scanf()</i> . The formatting string determines what additional arguments you need to provide. |
| <i>arg</i>       | A variable-argument list of the additional arguments, which you must have initialized with the <i>va_start()</i> macro.                                                     |

## Library:

`libc`

## Description:

The *vsscanf()* function scans input from the string designated by *in\_string*, under control of the argument *format*.

The *vsscanf()* function is a “varargs” version of *sscanf()*.

## Returns:

The number of input arguments for which values were successfully scanned and stored is returned, or EOF when the scanning is terminated by reaching the end of the input string.

**Examples:**

```
#include <stdio.h>
#include <stdarg.h>
#include <stdlib.h>

void sfind(char* string, char* format, ...)
{
 va_list arglist;

 va_start(arglist, format);
 vsscanf(string, format, arglist);
 va_end(arglist);
}

int main(void)
{
 int day, year;
 char weekday[10], month[12];

 sfind("Monday June 28 1999",
 "%s %s %d %d",
 weekday, month, &day, &year);
 printf("\n%s, %s %d, %d\n",
 weekday, month, day, year);
 return EXIT_SUCCESS;
}
```

**Classification:**

Unix

**Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Yes                     |



**Caveats:**

It's safe to call *vsscanf()* in a signal handler if the data isn't floating point.

**See also:**

*fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*, *vfprintf()*, *vfwscanf()*, *vscanf()*, *vswscanf()*, *vwscanf()*, *wscanf()*

## ***vswprintf()***

© 2004, QNX Software Systems Ltd.

*Write wide-character formatted output to a buffer (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vswprintf(wchar_t * buf,
 size_t n,
 const wchar_t * format,
 va_list arg);
```

### **Arguments:**

- |               |                                                                                                                                                                                              |
|---------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>buf</i>    | A pointer to the buffer where you want to function to store the formatted string.                                                                                                            |
| <i>n</i>      | The maximum number of wide characters to store in the buffer, including a terminating null character.                                                                                        |
| <i>format</i> | A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see <i>printf()</i> . |
| <i>arg</i>    | A variable-argument list of the additional arguments, which you must have initialized with the <i>va_start()</i> macro.                                                                      |

### **Library:**

`libc`

### **Description:**

The *vswprintf()* function formats data under control of the *format* control string, and writes the result to *buf*.

The *vswprint()* function is the wide-character version of *vsprintf()*, and is a “varargs” version of *swprintf()*.

## Returns:

The number of wide characters written, excluding the terminating **NUL**, or a negative number if an error occurred (*errno* is set).

## Classification:

ANSI

### Safety

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Yes                     |

## Caveats:

It's safe to call *vswprintf()* in a signal handler if the data isn't floating point.

## See also:

*fprintf()*, *fwprintf()*, *printf()*, *snprintf()*, *sprintf()*, *swprintf()*, *va\_start()*, *vfprintf()*, *vfwprintf()*, *vprintf()*, *vsnprintf()*, *vsprintf()*, *vwprintf()*, *wprintf()*

## ***vswscanf()***

© 2004, QNX Software Systems Ltd.

*Scan input from a wide-character string (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vswscanf(const wchar_t * ws,
 const wchar_t * format,
 va_list arg);
```

### **Arguments:**

*ws*            The wide-character string that you want to read from.

*format*        A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

*arg*            A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vswscanf()* function scans input from the string designated by *ws*, under control of the argument *format*.

The *vswscanf()* function is the wide-character version of *vscanf()*, and is a “varargs” version of *swscanf()*.

### **Returns:**

The number of input arguments for which values were successfully scanned and stored is returned, or EOF when the scanning is terminated by reaching the end of the input string.

## **Classification:**

ANSI

### **Safety**

---

|                    |                         |
|--------------------|-------------------------|
| Cancellation point | No                      |
| Interrupt handler  | No                      |
| Signal handler     | Read the <i>Caveats</i> |
| Thread             | Yes                     |

## **Caveats:**

It's safe to call *vswscanf()* in a signal handler if the data isn't floating point.

## **See also:**

*fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*, *vfscanf()*, *vwscanf()*, *vscanf()*, *vsscanf()*, *vswscanf()*, *wscanf()*

# ***vsyslog()***

© 2004, QNX Software Systems Ltd.

*Control system log (varargs)*

## **Synopsis:**

```
#include <syslog.h>
#include <stdarg.h>

void vsyslog(int priority,
 const char *message,
 va_list args);
```

## **Arguments:**

- |                 |                                                                                                                                                                                                                                                                      |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>priority</i> | The priority of the message; see “Message levels,” in the documentation for <i>syslog()</i> .                                                                                                                                                                        |
| <i>message</i>  | The message that you want to write. This message is identical to a <i>printf()</i> -format string, except that <i>%m</i> is replaced by the current error message (as denoted by the global variable <i>errno</i> ). A trailing newline is added if none is present. |
| <i>args</i>     | A variable-argument list of the additional arguments, which you must have initialized with the <i>va_start()</i> macro.                                                                                                                                              |

## **Library:**

`libc`

## **Description:**

The *vsyslog()* function writes *message* to the system message logger. The message is then written to the system console, log files, and logged-in users, or forwarded to other machines as appropriate. (See the `syslogd` command.)

This function is a “varargs” version of *syslog()*.

## **Classification:**

Unix

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | No  |

## **See also:**

*closelog(), openlog(), setlogmask(), syslog()*  
*logger, syslogd* in the *Utilities Reference*

## ***vwarn()*, *vwarnx()***

© 2004, QNX Software Systems Ltd.

*Formatted error message (varargs)*

### **Synopsis:**

```
#include <err.h>

void vwarn(const char *fmt,
 va_list args);

void vwarnx(const char *fmt,
 va_list args);
```

### **Arguments:**

*fmt*      NULL, or a *printf()*-style string used to format the message.

*args*     A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *err()* and *warn()* family of functions display a formatted error message on *stderr*. For a comparison of the members of this family, see *err()*.

The *vwarn()* function produces a message that consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, followed by a colon and a space, if the *fmt* argument isn't NULL
- the string associated with the current value of *errno*
- a newline character.

The *vwarnx()* function produces a similar message, except that it doesn't include the string associated with *errno*. The message consists of:



- the last component of the program name, followed by a colon and a space
- the formatted message, if the *fmt* argument isn't NULL
- a newline character.

### **Classification:**

Unix

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### **See also:**

*err()*, *errx()*, *stderr*, *strerror()*, *verr()*, *verrx()*, *warn()*, *warnx()*

## ***vwprintf()***

© 2004, QNX Software Systems Ltd.

*Write formatted output to standard output (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vwprintf(const wchar_t * format,
 va_list arg);
```

### **Arguments:**

*format*     A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

*arg*        A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vwprintf()* function writes output to the file *stdout*, under control of the argument *format*.

The *vwprintf()* function is the wide-character version of *vprintf()*, and is a “varargs” version of *wprintf()*.

### **Returns:**

The number of characters written, or a negative value if an output error occurred (*errno* is set).

### **Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, fprintf(), fwprintf(), printf(), snprintf(), sprintf(), swprintf(), va\_start(), vfprintf(), vfwprintf(), vprintf(), vsnprintf(), vsprintf(), vswprintf(), wprintf()*

## ***vwscanf()***

© 2004, QNX Software Systems Ltd.

*Scan wide-character input from standard input (varargs)*

### **Synopsis:**

```
#include <wchar.h>
#include <stdarg.h>

int vwscanf(const wchar_t * format,
 va_list arg);
```

### **Arguments:**

- format*     A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.
- arg*        A variable-argument list of the additional arguments, which you must have initialized with the *va\_start()* macro.

### **Library:**

`libc`

### **Description:**

The *vwscanf()* function scans input from the file designated by *stdin*, under control of the argument *format*.

The *vwscanf()* function is the wide-character version of *vscanf()*, and is a “varargs” version of *wscanf()*.

### **Returns:**

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning reached the end of the input stream before storing any values.

## **Classification:**

ANSI

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

## **See also:**

*fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *va\_start()*, *vfscanf()*,  
*vwscanf()*, *vscanf()*, *vsscanf()*, *vwscanf()*, *wscanf()*

## ***wait()***

© 2004, QNX Software Systems Ltd.

*Wait for the status of a terminated child process*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/wait.h>

pid_t wait(int * stat_loc);
```

### **Arguments:**

*stat\_loc*     NULL, or a pointer to a location where the function can store the terminating status of the child. For more information, see “Status macros,” below.

### **Library:**

`libc`

### **Description:**

The *wait()* function suspends execution of the calling process until status information from one of its terminated child processes is available, or until the delivery of a signal whose action is either to terminate the process or execute a signal handler. If status information is available prior to the call to *wait()*, the return is immediate.

### **Status macros**

If the *stat\_loc* variable is non-NULL, the terminating status of the child process is in the location that it points to. The macros listed below, defined in `<sys/wait.h>`, extract information from *stat\_loc*. The *stat\_val* argument to these macros is the integer value pointed to by *stat\_loc*.

POSIX defines the following macros:

*WEXITSTATUS( stat\_val )*

Evaluates to the low-order 8 bits of the termination status of the child process if the value of *WIFEXITED( stat\_val )* is nonzero.

*WIFCONTINUED( stat\_val )*

Evaluates to a nonzero value if the status returned was from a child process that has continued from a job control stop.

*WIFEXITED( stat\_val )*

Evaluates to a nonzero value if the status returned was from a normally terminated child process.

*WIFSIGNALED( stat\_val )*

Evaluates to nonzero value if the child process terminated from reception of a signal that wasn't caught.

*WIFSTOPPED( stat\_val )*

Evaluates to a nonzero value if the status returned is for a child process that's stopped.

*WSTOPSIG( stat\_val )*

Evaluates to the number of the signal that caused the child process to stop if the value of *WIFSTOPPED( stat\_val )* is nonzero.

*WTERMSIG( stat\_val )*

Evaluates to the number of the signal that terminated the child process if the value of *WIFSIGNALED( stat\_val )* is nonzero.

This macro isn't part of a POSIX standard:

*WCOREDUMP( stat\_val )*

Evaluates to a nonzero value if the child process left a core dump.

One of the macros *WIFEXITED(\*stat\_loc)* and *WIFSIGNALED(\*stat\_loc)* evaluates to a nonzero value.

The non-POSIX *waitid()* function gives even more status information than the above macros.

**Returns:**

The process ID of the terminating child process, or -1 if an error occurred or on delivery of a signal (*errno* is set to EINTR).

**Errors:**

- ECHILD      The calling process has no existing unwaited-for child processes.
- EINTR        The function was interrupted by a signal. The value of the location pointed to by *stat\_loc* is undefined.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*, *spawn()*, *wait4()*, *waitid()*, *waitpid()*



**Synopsis:**

```
#include <sys/wait.h>

pid_t wait3(int * stat_loc
 int options,
 struct rusage * resource_usage);
```

**Arguments:**

*stat\_loc*     NULL, or a pointer a location where the function can store the terminating status of the child process. For information about macros that extract information from this status, see “Status macros” in the documentation for *wait()*.

*options*     A combination of zero or more of the following flags:

- WCONTINUED — return the status for any child that was stopped and has been continued.
- WEXITED — wait for the process(es) to exit.
- WNOHANG — return immediately if there are no children to wait for.
- WNOWAIT — keep the process in a waitable state. This doesn't affect the state of the process; the process may be waited for again after this call completion.
- WSTOPPED — wait for and return the process status of any child that has stopped because it received a signal.
- WUNTRACED — report the status of a stopped child process.

*resource\_usage*

NULL, or a pointer to a **rusage** structure where the function can store information about resource usage. For information about this structure, see *getrusage()*.

**Library:**

libc

**Description:**

The *wait3()* function allows the calling thread to obtain status information for specified child processes.

The following call:

```
wait3(stat_loc, options, resource_usage);
```

is equivalent to the call:

```
waitpid((pid_t)-1, stat_loc, options);
```

except that on successful completion, if the *resource\_usage* argument to *wait3()* isn't a null pointer, the **rusage** structure that the third argument points to is filled in for the child process identified by the return value.

It's also equivalent to:

```
wait4((pid_t)-1, stat_loc, options, resource_usage);
```

**Returns:**

If the status of a child process is available, a value equal to the process ID of the child process for which status is reported.

If a signal is delivered to the calling process, -1 and *errno* is set to EINTR.

Zero if *wait3()* is invoked with WNOHANG set in *options* and at least one child process is specified by *pid* for which status isn't available, and status isn't available for any process specified by *pid*.

Otherwise, **(pid\_t)-1** and *errno* is set.

## Errors:

ECHILD      The calling process has no existing unwaited-for child processes, or the set of processes specified by the argument *pid* can never be in the states specified by the argument options.

## Classification:

Standard unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## Caveats:

New applications should use *waitpid()*.

## See also:

*exit()*, *fork()*, *pause()*, *wait4()*, *waitid()*, *waitpid()*

## **wait4()**

© 2004, QNX Software Systems Ltd.

*Wait for one or more child process to change its state*

### **Synopsis:**

```
#include <sys/types.h>
#include <sys/wait.h>

pid_t wait4(pid_t pid,
 int * stat_loc,
 int options,
 struct rusage * resource_usage);
```

### **Arguments:**

- pid*            The set of child processes that you want to get status information for:
- less than -1 — any child process whose process group ID is equal to the absolute value of *pid*.
  - -1 — any child process
  - 0 — any child process whose process group ID is equal to that of the calling process.
  - greater than 0 — the single child process with this ID.
- stat\_loc*        NULL, or a pointer a location where the function can store the terminating status of the child process. For information about macros that extract information from this status, see “Status macros” in the documentation for *wait()*.
- options*        A combination of zero or more of the following flags:
- WCONTINUED — return the status for any child that was stopped and has been continued.
  - WEXITED — wait for the process(es) to exit.
  - WNOHANG — return immediately if there are no children to wait for.
  - WNOWAIT — keep the process in a waitable state. This doesn't affect the state of the process; the process may be waited for again after this call completion.

- WSTOPPED — wait for and return the process status of any child that has stopped because it received a signal.
- WUNTRACED — report the status of a stopped child process.

*resource\_usage*

NULL, or a pointer to a **rusage** structure where the function can store information about resource usage. For information about this structure, see *getrusage()*.

**Library:**

`libc`

**Description:**

The *wait4()* function suspends execution of the calling process until status information from one of its terminated child processes is available, or until the delivery of a signal whose action is either to terminate the process or execute a signal handler. If status information is available prior to the call to *wait4()*, the return is immediate.

The *wait4()* function behaves the same as the *wait()* function when passed a *pid* argument of -1, and the *options* argument has a value of zero.

Only one of the *WIFEXITED(stat\_val)* and *WIFSIGNALED(stat\_val)* macros can evaluate to a nonzero value.

The following call:

```
wait3(stat_loc, options, resource_usage);
```

is equivalent to the call:

```
waitpid((pid_t)-1, stat_loc, options);
```

except that on successful completion, if the *resource\_usage* argument to *wait3()* isn't a NULL pointer, the **rusage** structure that the third

argument points to is filled in for the child process identified by the return value.

It's also equivalent to:

```
wait4((pid_t)-1, stat_loc, options, resource_usage);
```

## Returns:

If successful, *wait4()* returns the process id of the terminating child process. If *wait4()* was invoked with WNOHANG set in *options*, it has at least one child process specified by *pid* for which status is not available, and status is not available for any process specified by *pid*, a value of zero is returned. On delivery of a signal *waitpid()* returns -1, and *errno* is set to EINTR.

## Errors:

|        |                                                                                                                 |
|--------|-----------------------------------------------------------------------------------------------------------------|
| ECHILD | The calling process has no existing unwaited-for child processes that meet the criteria set by <i>pid</i> .     |
| EINTR  | The function was interrupted by a signal. The value of the location pointed to by <i>stat_loc</i> is undefined. |
| EINVAL | The value of the <i>options</i> argument isn't valid.                                                           |

## Classification:

Unix

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*exit(), fork(), pause(), wait(), wait3(), waitid(), waitpid()*

## ***waitid()***

© 2004, QNX Software Systems Ltd.

*Wait for a child process to change state*

### **Synopsis:**

```
#include <sys/wait.h>

int waitid(idtype_t idtype,
 id_t id,
 siginfo_t * infop,
 int options);
```

### **Arguments:**

- idtype* Which children you want to wait for:
- P\_PID — the child with a process ID of (*pid\_t*) *id*.
  - P\_PGID — any child with a process group ID equal to (*pid\_t*) *id*.
  - P\_ALL — any child; *id* is ignored.
- id* The process or process group ID that you want to wait for, depending on the value of *idtype*.
- infop* A pointer to a **siginfo\_t** structure, as defined in **<sys/siginfo.h>**, where the function can store the current state of the child; see below.
- options* A combination of zero or more of the following flags:
- WCONTINUED — return the status for any child that was stopped and has been continued.
  - WEXITED — wait for the process(es) to exit.
  - WNOHANG — return immediately if there are no children to wait for.
  - WNOWAIT — keep the process in a waitable state. This doesn't affect the state of the process; the process may be waited for again after this call completion.
  - WSTOPPED — wait for and return the process status of any child that has stopped because it received a signal.
  - WUNTRACED — report the status of a stopped child process.



## Library:

libc

## Description:

The *waitid()* function suspends the calling process until one of its children changes state. It records the current state of a child in the structure pointed to by *infop*. If a child process changed state prior to the call to *waitid()*, *waitid()* returns immediately.

If *waitid()* returns because a child process was found that satisfied the conditions indicated by the arguments *idtype* and *options*, then the structure pointed to by *infop* is filled in by the system with the status of the process. The *si\_signo* member is always SIGCHLD.

If *idtype* is P\_ALL and *options* is **WEXITED** | **WTRAPPED**, *waitid()* is equivalent to *wait()*.

## Returns:

- 0 One of the children changed its state. If WNOHANG was used, 0 can be returned (indicating no error); however, no children may have changed state if *info->si\_pid* is 0.
- 1 An error occurred (*errno* is set).

## Errors:

- ECHILD The set of processes specified by *idtype* and *id* doesn't contain any unwaited-for processes.
- EFAULT The *infop* argument points to an illegal address.
- EINTR The *waitid()* function was interrupted by a signal.
- EINVAL An invalid value was specified for *options*, or *idtype* and *id* specify an invalid set of processes.

**Classification:**

Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*execl(), execlp(), execlpe(), execv(), execvp(),  
execvpe(), exit(), fork(), pause(), sigaction(), signal(), wait(), wait3(),  
wait4(), waitpid()*

**Synopsis:**

```
#include <sys/types.h>
#include <sys/wait.h>

pid_t waitpid(pid_t pid,
 int * stat_loc,
 int options);
```

**Arguments:**

- pid*            The set of child processes that you want to get status information for:
- less than -1 — any child process whose process group ID is equal to the absolute value of *pid*.
  - -1 — any child process
  - 0 — any child process whose process group ID is equal to that of the calling process.
  - greater than 0 — the single child process with this ID.
- stat\_loc*        NULL, or a pointer a location where the function can store the terminating status of the child process. For information about macros that extract information from this status, see “Status macros” in the documentation for *wait()*.
- options*        A combination of zero or more of the following flags:
- WCONTINUED — return the status for any child that was stopped and has been continued.
  - WEXITED — wait for the process(es) to exit.
  - WNOHANG — return immediately if there are no children to wait for.
  - WNOWAIT — keep the process in a waitable state. This doesn't affect the state of the process; the process may be waited for again after this call completion.

- **WSTOPPED** — wait for and return the process status of any child that has stopped because it received a signal.
- **WUNTRACED** — report the status of a stopped child process.

**Library:**

`libc`

**Description:**

The *waitpid()* function suspends execution of the calling process until status information from one of its terminated child processes is available, or until the delivery of a signal whose action is either to terminate the process or execute a signal handler. If status information is available prior to the call to *waitpid()*, the return is immediate.

The *waitpid()* function behaves the same as *wait()* when passed a *pid* argument of -1, and the *options* argument has a value of zero.

Only one of the *WIFEXITED(stat\_val)* and *WIFSIGNALED(stat\_val)* macros can evaluate to a nonzero value.

**Returns:**

The process ID of the terminating child process on success. If *waitpid()* is invoked with **WNOHANG** set in *options*, it has at least one child process specified by *pid* for which status isn't available, and status isn't available for any process specified by *pid*, a value of zero is returned. On delivery of a signal, *waitpid()* returns -1, and *errno* is set to **EINTR**.

**Errors:**

- |               |                                                                                                                 |
|---------------|-----------------------------------------------------------------------------------------------------------------|
| <b>ECHILD</b> | The calling process has no existing unwaited-for child processes that meet the criteria set by <i>pid</i> .     |
| <b>EINTR</b>  | The function was interrupted by a signal. The value of the location pointed to by <i>stat_loc</i> is undefined. |

EINVAL     The value of the *options* argument isn't valid.

**Classification:**

POSIX 1003.1

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*spawn()*, *wait()*, *wait3()*, *wait4()*, *waitid()*

## ***warn()*, *warnx()***

© 2004, QNX Software Systems Ltd.

*Formatted error message*

---

### **Synopsis:**

```
#include <err.h>

void warn(const char* fmt, ...);

void warnx(const char* fmt, ...);
```

### **Arguments:**

*fmt* NULL, or a *printf()*-style string used to format the message.

Additional arguments

As required by the format string.

### **Library:**

`libc`

### **Description:**

The *err()* and *warn()* family of functions display a formatted error message on *stderr*. For a comparison of the members of this family, see *err()*.

The *warn()* function produces a message that consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, followed by a colon and a space, if the *fmt* argument isn't NULL
- the string associated with the current value of *errno*
- a newline character.

The *warnx()* function produces a similar message, except that it doesn't include the string associated with *errno*. The message consists of:

- the last component of the program name, followed by a colon and a space
- the formatted message, if the *fmt* argument isn't NULL
- a newline character.

## Examples:

Warn of an error:

```
if ((fd = open(raw_device, O_RDONLY, 0)) == -1)
 warnx("%s: %s: trying the block device",
 raw_device, strerror(errno));
if ((fd = open(block_device, O_RDONLY, 0)) == -1)
 warn("%s", block_device);
```

## Classification:

Unix

### Safety

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*err()*, *errx()*, *stderr*, *strerror()*, *verr()*, *verrx()*, *vwarn()*, *vwarnx()*

## **wcrtomb()**

© 2004, QNX Software Systems Ltd.

*Convert a wide-character code to a character*

### **Synopsis:**

```
#include <wchar.h>

size_t wcrtomb(char * s,
 wchar_t wc,
 mbstate_t * ps);
```

### **Arguments:**

- s*      NULL, or a pointer to a location where the function can store the multibyte character.
- wc*     The wide character that you want to convert.
- ps*     An internal pointer that lets *wcrtomb()* be a restartable version of *wctomb()*; if *ps* is NULL, *wcrtomb()* uses its own internal variable.  
You can call *mbstate\_t* to determine the status of this variable.

### **Library:**

`libc`

### **Description:**

The *wcrtomb()* function determines the number of bytes needed to represent the wide character *wc* as a multibyte character and stores the multibyte character in the location pointed to by *s*, to a maximum of MB\_CUR\_MAX bytes.

This function is affected by LC\_CTYPE.

### **Returns:**

The number of bytes stored, or `(size_t) -1` if the variable *wc* is an invalid wide-character code.



## Errors:

- EILSEQ Invalid wide-character code.
- EINVAL The variable *ps* points to an invalid conversion state.

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

# **wcscat()**

© 2004, QNX Software Systems Ltd.

*Concatenate two wide-character strings*

## **Synopsis:**

```
#include <wchar.h>

wchar_t * wcscat(wchar_t * ws1,
 const wchar_t * ws2);
```

## **Arguments:**

*ws1, ws2*     The wide-character strings that you want to concatenate.

## **Library:**

`libc`

## **Description:**

The *wcscat()* function appends a copy of the string pointed to by *ws2*, including the terminating NUL wide character, to the end of the string pointed to by *ws1*. The first wide character of *ws2* overwrites the NUL wide character at the end of *ws1*.

## **Returns:**

The same pointer as *ws1*.

## **Classification:**

ANSI

### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## **wcschr()**

© 2004, QNX Software Systems Ltd.

*Find the first occurrence of a wide character in a string*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcschr(const wchar_t * ws,
 wchar_t wc);
```

### **Arguments:**

*ws*     The wide-character string that you want to search.

*wc*     The wide character that you're looking for.

### **Library:**

`libc`

### **Description:**

The *wcschr()* function finds the first occurrence of *wc* in the string pointed to by *ws*. The terminating NUL character is considered to be part of the string.

### **Returns:**

A pointer to the located wide character, or NULL if *wc* doesn't occur in the string.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*,  
*wcsstr()*, *wcstok()*

## **wcscmp()**

© 2004, QNX Software Systems Ltd.

*Compare two wide-character strings*

### **Synopsis:**

```
#include <wchar.h>

int wcscmp(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*, *ws2*     The wide-character strings that you want to compare.

### **Library:**

`libc`

### **Description:**

The *wcscmp()* function compares the wide-character strings pointed to by *ws1* and *ws2*.

### **Returns:**

< 0     *ws1* is less than *ws2*.  
0        *ws1* is equal to *ws2*.  
> 0     *ws1* is greater than *ws2*.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strcasecmp(), strcmp(), strcmpi(), strcoll(), stricmp(), strncasecmp(),  
strncmp(), strnicmp(), wcscoll(), wcsncmp()*

## ***wcscoll()***

© 2004, QNX Software Systems Ltd.

*Compare two wide-character strings, using the locale's collating sequence*

### **Synopsis:**

```
#include <wchar.h>

int wcscoll(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*, *ws2*     The wide-character strings that you want to compare.

### **Library:**

`libc`

### **Description:**

The *wcscoll()* function compares the wide-character strings pointed to by *ws1* and *ws2*, using the LC\_COLLATE collating sequence selected by the *setlocale()* function.

### **Returns:**

< 0     *ws1* is less than *ws2*.  
0       *ws1* is equal to *ws2*.  
> 0     *ws1* is greater than *ws2*.

### **Classification:**

ANSI

### ***wcscoll()***

#### **Safety**

---

Cancellation point    No

Interrupt handler     No

*continued...*



**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | No  |

**See also:**

*setlocale(), strcasecmp(), strcmp(), strcmpi(), strcoll(), stricmp(), strncasecmp(), strncmp(), strnicmp(), wscmp(), wcsncmp()*

## **wcscpy()**

© 2004, QNX Software Systems Ltd.

*Copy a wide-character string*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcscpy(wchar_t * ws1,
 const char * ws2);
```

### **Arguments:**

*ws1*     A pointer to where you want to copy the string.  
*ws2*     The wide-character string that you want to copy.

### **Library:**

`libc`

### **Description:**

The *wcscpy()* function copies the string pointed to by *ws2*, including the terminating NUL wide character, into the array pointed to by *ws1*.



---

This function isn't guaranteed to work properly for copying overlapping strings; use *wmemmove()* instead.

---

### **Returns:**

The same pointer as *ws1*.

### **Classification:**

ANSI

#### **Safety**

---

Cancellation point   No  
Interrupt handler     Yes

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*memmove(), strcpy(), strdup(), strncpy(), wcsncpy(), wmemmove()*

## ***wcscspn()***

© 2004, QNX Software Systems Ltd.

*Count the wide characters at the beginning of a string that aren't in a given character set*

### **Synopsis:**

```
#include <wchar.h>

size_t wcscspn(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*     The wide-character string that you want to search.  
*ws2*     The set of wide characters you want to look for.

### **Library:**

`libc`

### **Description:**

The *strspn()* function returns the length of the initial segment of the string pointed to by *ws1* consisting entirely of wide characters *not* from the string pointed to by *ws2*. The terminating NUL isn't considered to be part of *ws2*.

### **Returns:**

The length of the segment.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcschr()*, *wcspbrk()*, *wcsrchr()*, *wcsspn()*, *wcsstr()*,  
*wcstok()*

## ***wcsftime()***

© 2004, QNX Software Systems Ltd.

*Format the time into a wide-character string*

### **Synopsis:**

```
#include <wchar.h>

size_t wcsftime(wchar_t * wcs,
 size_t maxsize,
 const wchar_t * format,
 const struct tm * timeptr);
```

### **Arguments:**

|                |                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------|
| <i>wcs</i>     | A pointer to a buffer where the function can store the formatted time.                                 |
| <i>maxsize</i> | The maximum size of the buffer.                                                                        |
| <i>format</i>  | The format that you want to use for the time; see “Formats,” in the description of <i>strftime()</i> . |
| <i>timeptr</i> | A pointer to a <code>tm</code> structure that contains the time that you want to format.               |

### **Library:**

`libc`

### **Description:**

The *wcsftime()* function is similar to *strftime()*, except that *wcsftime()* works with wide characters.

### **Returns:**

The number of wide characters placed into the array, not including the terminating null character, or 0 if the number of wide characters exceeds *maxsize* (in this case, the string contents are indeterminate).

If an error occurs, *errno* is set.

## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*asctime()*, *asctime\_r()*, *ctime()*, *ctime\_r()*, *sprintf()*, *strftime()*, **tm**, *tzset()*

## **wcslen()**

© 2004, QNX Software Systems Ltd.

*Compute the length of a wide-character string*

### **Synopsis:**

```
#include <wchar.h>

size_t wcslen(const wchar_t * ws);
```

### **Arguments:**

*ws* The wide-character string whose length you want to calculate.

### **Library:**

`libc`

### **Description:**

The *wcslen()* function counts the wide characters in the string pointed to by *ws*.

### **Returns:**

The number of wide characters, not counting the terminating NUL.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*strlen()*

## **wcsncat()**

© 2004, QNX Software Systems Ltd.

*Concatenate two wide-character strings, up to a maximum length*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcsncat(wchar_t * ws1,
 const wchar_t * ws2
 size_t n);
```

### **Arguments:**

*ws1*, *ws2*     The wide-character strings that you want to concatenate.

*n*             The maximum number of wide characters that you want to add from the *ws2* string.

### **Library:**

`libc`

### **Description:**

The *wcsncat()* function appends a copy of the string pointed to by *ws2*, including the terminating NUL wide character, to the end of the string pointed to by *ws1*. The first character of *ws2* overwrites the NUL wide character at the end of *ws1*. The function writes no more than *n* wide characters from *ws2* and appends a NUL wide character to the result.

### **Returns:**

The same pointer as *ws1*.

### **Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## ***wcsncmp()***

© 2004, QNX Software Systems Ltd.

*Compare two wide-character strings, up to a given length*

### **Synopsis:**

```
#include <wchar.h>

int wcsncmp(const wchar_t * ws1,
 const wchar_t * ws2,
 size_t n);
```

### **Arguments:**

*ws1, ws2*     The wide-character strings that you want to compare.

*n*             The maximum number of wide characters that you want to compare.

### **Library:**

`libc`

### **Description:**

The *wcsncmp()* function compares up to *n* wide characters from the strings pointed to by *ws1* and *ws2*.

### **Returns:**

< 0     *ws1* is less than *ws2*.

0        *ws1* is equal to *ws2*.

> 0     *ws1* is greater than *ws2*.

### **Classification:**

ANSI

---

#### **Safety**

Cancellation point   No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*strcasecmp()*, *strcmp()*, *strcmpi()*, *strcoll()*, *stricmp()*, *strncasecmp()*,  
*strncmp()*, *strnicmp()*, *wscmp()*, *wscoll()*

## ***wcsncpy()***

© 2004, QNX Software Systems Ltd.

*Copy a wide-character string, to a maximum length*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcsncpy(wchar_t * ws1,
 const char * ws2,
 size_t n);
```

### **Arguments:**

- ws1*     A pointer to where you want to copy the wide-character string.
- ws2*     The wide-character string that you want to copy.
- n*        The maximum number of wide characters that you want to copy.

### **Library:**

`libc`

### **Description:**

The *wcsncpy()* function copies the string pointed to by *ws2*, including the terminating NUL wide character, into the array pointed to by *ws1*, to a maximum of *n* wide characters. It adds NUL characters if *ws2* has fewer than *n* characters but doesn't add a NUL if *ws2* has more.



This function isn't guaranteed to work properly for copying overlapping strings; use *wmemmove()* instead.

---

### **Returns:**

The same pointer as *ws1*.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memmove()*, *strcpy()*, *strdup()*, *strncpy()*, *wscpy()*, *wmemmove()*

## **wcspbrk()**

© 2004, QNX Software Systems Ltd.

*Find the first wide character in a string that's in a given character set*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcspbrk(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*     The wide-character string that you want to search.  
*ws2*     The set of wide characters you want to look for.

### **Library:**

`libc`

### **Description:**

The *wcspbrk()* function locates the first occurrence in the string pointed to by *ws1* of *any* wide character from the string pointed to by *ws2*.

### **Returns:**

A pointer to the located character, or NULL if no character from *ws2* occurs in *ws1*.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcschr()*, *wcscspn()*, *wcsrchr()*, *wcsspn()*, *wcsstr()*,  
*wcstok()*

## **wcsrchr()**

© 2004, QNX Software Systems Ltd.

*Find the last occurrence of a wide character in a string*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcsrchr(const wchar_t * ws,
 wchar_t wc);
```

### **Arguments:**

*ws*     The wide-character string that you want to search.

*wc*     The wide character that you're looking for.

### **Library:**

`libc`

### **Description:**

The *wcsrchr()* function finds the last occurrence of *wc* in the string pointed to by *ws*. The terminating NUL character is considered to be part of the string.

### **Returns:**

A pointer to the located wide character, or NULL if *wc* doesn't occur in the string.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsspn()*,  
*wcsstr()*, *wcstok()*

## ***wcsrtombs()***

© 2004, QNX Software Systems Ltd.

*Convert a wide-character string into a multibyte character string (restartable)*

### **Synopsis:**

```
#include <wchar.h>

size_t wcsrtombs(char * dst,
 const wchar_t ** src,
 size_t len,
 mbstate_t * ps);
```

### **Arguments:**

- dst*     A pointer to a buffer where the function can store the multibyte-character string.
- src*     A pointer to the wide-character string that you want to convert.
- len*     The maximum number of multibyte characters to store.
- ps*     An internal pointer that lets *wcsrtombs()* be a restartable version of *wcstombs()*; if *ps* is NULL, *wcsrtombs()* uses its own internal variable.
- You can call *mbstate\_t* to determine the status of this variable.

### **Library:**

`libc`

### **Description:**

The *wcsrtombs()* function converts a string of wide-characters pointed to by *src* into the corresponding multi-byte characters pointed to by *dst*. No more than *len* bytes are stored, including the terminating NULL character.

The function converts each character as if by a call to *wctomb()* and stops early if:

- A sequence of bytes doesn't conform to a valid character.

- Converting the next character would exceed the limit of *len* total bytes.

The *wcsrtombs()* function uses *ps* to make it thread safe; if *ps* is a NULL pointer, *wcsrtombs()* uses its own internal pointer.

### Returns:

The number of total bytes successfully converted, not including the terminating NULL byte, or `(size_t) - 1` if an invalid wide-character code was found.

### Classification:

ANSI

#### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

### See also:

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## **wcsspn()**

© 2004, QNX Software Systems Ltd.

*Count the wide characters at the beginning of a string that are in a given character set*

### **Synopsis:**

```
#include <wchar.h>

size_t wcsspn(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*     The wide-character string that you want to search.  
*ws2*     The set of wide characters you want to look for.

### **Library:**

`libc`

### **Description:**

The *wcsspn()* function returns the length of the initial segment of the string pointed to by *ws1* consisting entirely of wide characters from the string pointed to by *ws2*. The terminating NUL isn't considered to be part of *ws2*.

### **Returns:**

The length of the segment.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*,  
*wcsstr()*, *wcstok()*

## **wcsstr()**

© 2004, QNX Software Systems Ltd.

*Find one wide-character string inside another*

### **Synopsis:**

```
#include <wchar.h>

wchar_t * wcsstr(const wchar_t * ws1,
 const wchar_t * ws2);
```

### **Arguments:**

*ws1*     The wide-character string that you want to search.  
*ws2*     The wide-character string that you're looking for.

### **Library:**

`libc`

### **Description:**

The *wcsstr()* function locates the first occurrence in the string pointed to by *ws1* of the sequence of wide characters, excluding the terminating NUL, in the string pointed to by *ws2*.

### **Returns:**

A pointer to the located string, NULL if the string wasn't found, or the same pointer as *ws1* if *ws2* points to a zero-length string.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |



**See also:**

*memchr()*, *strchr()*, *strcspn()*, *strpbrk()*, *strrchr()*, *strspn()*, *strstr()*,  
*strtok()*, *strtok\_r()*, *wcschr()*, *wcscspn()*, *wcspbrk()*, *wcsrchr()*,  
*wcsspn()*, *wcstok()*

## ***wcstod()*, *wcstof()*, *wcstold()***

© 2004, QNX Software Systems Ltd.

*Convert a wide-character string into a double, float, or long double*

### **Synopsis:**

```
#include <wchar.h>

double wcstod(const wchar_t * ptr,
 wchar_t ** endptr);

float wcstof(const wchar_t * ptr,
 wchar_t ** endptr);

long double wcstold(const wchar_t * ptr,
 wchar_t ** endptr);
```

### **Arguments:**

*nptr*        A pointer to the string to parse.

*endptr*     If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.

### **Library:**

`libc`

### **Description:**

These functions convert a wide-character string to a number:

- *wcstod()* function converts it to a **double**
- *wcstof()* converts it to a **float**
- *wcstold()* to a **long double**.

These functions recognize strings containing the following:

- optional white space
- an optional plus or minus sign
- a sequence of digits containing an optional decimal point

- an optional **e** or **E**, followed by an optionally signed sequence of digits.

The functions expect the string to have a plus or minus sign, followed by one of these forms:

- A sequence of decimal digits, optionally followed by a radix character, optionally followed by an exponent part.
- A **0x** or **0X** followed by a sequence of hexadecimal digits, optionally followed by a radix character, optionally followed by a binary exponent part.
- The case-insensitive string **INF** or **INFINITY**.
- The case-insensitive string **NAN** or **NAN** (*n-wchar-sequence*) where *n-wchar-sequence* may be a digit, a nondigit, a *n-wchar-sequence* digit or a *n-wchar-sequence* nondigit.

The value is correctly rounded if the subject is hexadecimal and `FLT_RADIX` is 2.

The radix character is locale specific, depending upon `LC_NUMERIC`.

The conversion ends at the first unrecognized character. If *endptr* isn't `NULL`, a pointer to the unrecognized wide character is stored in the object *endptr* points to.



---

Because 0 is a valid return that is also used for an error, you should set *errno* to 0 before calling these functions, and check *errno* again afterward. These functions don't change *errno* on success.

---

## Returns:

The converted value. If the correct value would cause overflow, plus or minus `HUGE_VAL` is returned according to the sign, and *errno* is set to `ERANGE`. If the correct value would cause underflow, then zero is returned, and *errno* is set to `ERANGE`.

Zero is returned when the input string can't be converted. When an error occurs, *errno* indicates the error detected.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## ***wcstoimax()*, *wcstoumax()***

*Convert a wide-character string into an integer*

### **Synopsis:**

```
#include <inttypes.h>

intmax_t wcstoimax (const wchar_t * nptr,
 wchar_t ** endptr,
 int base);

uintmax_t wcstoumax (const wchar_t * nptr,
 wchar_t ** endptr,
 int base);
```

### **Arguments:**

- nptr*      A pointer to the string to parse.
- endptr*    If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*      The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.

**Library:**

`libc`

**Description:**

The *wcstoimax()* and *wcstoumax()* functions are the same as the *wcstol()*, *wcstoll()*, *wcstoul()*, and *wcstoull()* functions except that they return objects of type `intmax_t` and `uintmax_t`.

**Returns:**

The converted value.

If the correct value causes an overflow, (`INTMAX_MAX` | `UINTMAX_MAX` or `INTMAX_MIN`) is returned according to the sign and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EINVAL`.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*strtol()*, *wcstol()*, *wcstoul()*

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

**Synopsis:**

```
#include <wchar.h>

wchar_t * wcstok(wchar_t * ws1,
 const wchar_t * ws2,
 wchar_t ** ptr);
```

**Arguments:**

- ws1* NULL, or the wide-character string that you want to break into tokens; see below.
- ws2* A set of the wide characters that separate the tokens.
- ptr* The address of a pointer to a **wchar\_t** object, which the function can use to store information necessary for it to continue scanning the same string.

**Library:**

**libc**

**Description:**

The function *wcstok()* breaks the wide-character string pointed to by *ws1* into a sequence of tokens, each of which is delimited by a wide character from the string pointed to by *ws2*.

In the first call to *wcstok()*, *ws1* must point to a null-terminated string, *ws2* must point to a null-terminated string of separator wide characters, and *ptr* is ignored. The *wcstok()* function returns a pointer to the first wide character of the first token, writes a NUL wide character into *ws1* immediately following the returned token, and updates *ptr*.

In subsequent calls, *ws1* must be NULL, and *ptr* must be unchanged from the previous call so that subsequent calls will move through the string *ws1*, returning successive tokens until no tokens remain. The separator string *ws2* may differ from call to call. When no tokens remain in *ws1*, a NULL pointer is returned.

**Returns:**

A pointer to the token found, or NULL if no token was found.

**Classification:**

POSIX 1003.1

*wcstok()*

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memchr(), strchr(), strcspn(), strpbrk(), strrchr(), strset(), strspn(), strstr(), strtok(), strtok\_r(), wcschr(), wcsespn(), wcpbrk(), wcsrchr(), wcsspn(), wcsstr()*



### **Synopsis:**

```
#include <stdlib.h>

long wcstol(const wchar_t * ptr,
 wchar_t ** endptr,
 int base);

long long wcstoll(const wchar_t * ptr,
 wchar_t ** endptr,
 int base);
```

### **Arguments:**

- ptr*            A pointer to the string to parse.
- endptr*        If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*           The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.

## Library:

`libc`

## Description:

The *wcstol()* function converts the string pointed to by *ptr* into a **long**; *wcstoll()* converts the string into a **long long**.

These functions recognize strings that contain the following:

- optional white space
- an optional plus or minus sign
- a sequence of digits and letters.

The conversion ends at the first unrecognized wide character. If *endptr* isn't NULL, a pointer to the unrecognized wide character is stored in the object *endptr* points to.

## Returns:

The converted value.

If the correct value causes an overflow, `LONG_MAX | LONGLONG_MAX` or `LONG_MIN | LONGLONG_MIN` is returned according to the sign, and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EDOM`.

## Errors:

|                     |                                                                                 |
|---------------------|---------------------------------------------------------------------------------|
| <code>ERANGE</code> | The value is not representable                                                  |
| <code>EINVAL</code> | The value for <i>base</i> is not supported or no conversion could be performed. |

## Classification:

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*errno*

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## ***wcstombs()***

© 2004, QNX Software Systems Ltd.

*Convert a wide-character string into a multibyte character string*

### **Synopsis:**

```
#include <stdlib.h>

size_t wcstombs(char* s,
 const wchar_t* pwcs,
 size_t n);
```

### **Arguments:**

*s*            A pointer to a buffer where the function can store the multibyte-character string.

*pwcs*        The wide-character string that you want to convert.

*n*            The maximum number of bytes to store.

### **Library:**

`libc`

### **Description:**

The *wcstombs()* function converts a sequence of wide character codes from the array pointed to by *pwcs* into a sequence of multibyte characters, and stores them in the array pointed to by *s*. It stops if a multibyte character exceeds the limit of *n* total bytes, or if the NUL character is stored. At most *n* bytes of the array pointed to by *s* are modified.

The *wcsrtombs()* function is a restartable version of *wcstombs()*.

### **Returns:**

The number of array elements modified, not including the terminating zero code, if present, or `(size_t) - 1` if an invalid multibyte character is encountered.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

wchar_t wbuffer[] = {
 0x0073,
 0x0074,
 0x0072,
 0x0069,
 0x006e,
 0x0067,
 0x0000
};

int main(void)
{
 char mbsbuffer[50];
 int i, len;

 len = wcstombs(mbsbuffer, wbuffer, 50);
 if(len != -1) {
 for(i = 0; i < len; i++)
 printf("%4.4x", wbuffer[i]);
 printf("\n");
 mbsbuffer[len] = '\0';
 printf("%s(%d)\n", mbsbuffer, len);
 }
 return EXIT_SUCCESS;
}
```

produces the output:

```
/0073/0074/0072/0069/006e/0067
string(6)
```

**Classification:**

ANSI

**Safety**

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | Yes |
| Thread         | Yes |

**See also:**

*mblen(), mbtowc(), mbstowcs(), wcsrtombs(), wctomb()*

**Synopsis:**

```
#include <stdlib.h>

long wcstoul(const wchar_t * ptr,
 wchar_t ** endptr,
 int base);

long long wcstoull(const wchar_t * ptr,
 char** endptr,
 int base);
```

**Arguments:**

- ptr*            A pointer to the string to parse.
- endptr*        If this argument isn't NULL, the function stores in it a pointer to the first unrecognized character found in the string.
- base*           The base of the number being parsed:
- If *base* is zero, the first characters after the optional sign determine the base used for the conversion. If the first characters are **0x** or **0X** the digits are treated as hexadecimal. If the first character is **0**, the digits are treated as octal. Otherwise, the digits are treated as decimal.
  - If *base* isn't zero, it must have a value between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters **0x** or **0X** may optionally precede the sequence of letters and digits.

## Library:

`libc`

## Description:

These functions convert a wide-character string into a number:

- *wcstoul()* converts the string into an **unsigned long**
- *wcstoull()* converts it into a **unsigned long long**.

These functions recognize a string containing optional white space, followed by a sequence of digits and letters. The conversion ends at the first unrecognized character. A pointer to that character is stored in the object *endptr* points to, if *endptr* isn't NULL.

If *base* is zero, the first characters determine the base used for the conversion. If the first characters are `0x` or `0X` the digits are treated as hexadecimal. If the first character is `0`, the digits are treated as octal. Otherwise, the digits are treated as decimal.

If *base* isn't zero, it must have a value of between 2 and 36. The letters a-z and A-Z represent the values 10 through 35. Only those letters whose designated values are less than *base* are permitted. If the value of *base* is 16, the characters `0x` or `0X` may optionally precede the sequence of letters and digits.

## Returns:

The converted value.

If the correct value causes an overflow, `ULONG_MAX` | `ULONGLONG_MAX` is returned and *errno* is set to `ERANGE`. If *base* is out of range, zero is returned and *errno* is set to `EDOM`.

## Errors:

|                     |                                                                                 |
|---------------------|---------------------------------------------------------------------------------|
| <code>ERANGE</code> | The value is not representable                                                  |
| <code>EINVAL</code> | The value for <i>base</i> is not supported or no conversion could be performed. |



## Classification:

ANSI

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*errno*

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## **wcscxfrm()**

© 2004, QNX Software Systems Ltd.

*Transform one wide-character string into another, to a given length*

### **Synopsis:**

```
#include <wchar.h>

int wcscxfrm(wchar_t * ws1,
 const wchar_t * ws2,
 size_t n);
```

### **Arguments:**

*ws1*     The string that you want to transform.

*ws2*     The string that you want to place in *dst*.

*n*        The maximum number of characters to transform.

### **Library:**

`libc`

### **Description:**

The *wcscxfrm()* function transforms the string pointed to by *ws2* to the buffer pointed to by *ws1*, to a maximum of *n* wide-characters, including the terminating null. The two strings shouldn't overlap.

A call to *wcscmp()* returns the same result for two strings transformed by *wcscxfrm()* as *wcscoll()* would return for the original versions of the strings.



---

This function doesn't report errors in its returns; set *errno* to 0, call *wcscxfrm()*, and then check *errno* again.

---

### **Returns:**

The length of the transformed wide-character string. If this value is greater than *n*, the contents of *ws1* are indeterminate.

**Classification:**

ANSI

***wcscxfrm()***

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*setlocale(), strxfrm()*

## **wctob()**

© 2004, QNX Software Systems Ltd.

*Convert a wide character to a single-byte code*

### **Synopsis:**

```
#include <wchar.h>

int wctob(wint_t c);
```

### **Arguments:**

*c* The wide character that you want to convert.

### **Library:**

`libc`

### **Description:**

The *wctob()* function returns the single-byte representation of a wide character.

This function is affected by LC\_CTYPE.

### **Returns:**

The single-byte representation, or EOF if *c* isn't a valid single-byte character.

### **Classification:**

ANSI

#### **Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## **wctomb()**

© 2004, QNX Software Systems Ltd.

*Convert a wide character into a multibyte character*

### **Synopsis:**

```
#include <stdlib.h>
int wctomb(char * s,
 wchar_t wc);
```

### **Arguments:**

- s*      NULL, or a pointer to a location where the function can store the multibyte character.
- wc*     The wide character that you want to convert.

### **Library:**

`libc`

### **Description:**

The *wctomb()* function determines the number of bytes required to represent the multibyte character corresponding to the code contained in *wc*. If *s* isn't NULL, the multibyte character representation is stored in the array it points to. At most MB\_CUR\_MAX characters are stored.

### **Returns:**

- If *s* is NULL:
  - 0      The *wctomb()* function uses locale specific multibyte character encoding that's not state-dependent.
  - >0     The function is state-dependent.
- If *s* isn't NULL:
  - 1     If the value of *wchar* doesn't correspond to a valid multibyte character.
  - x*     The number of bytes that comprise the multibyte character corresponding to the value of *wchar*.

**Examples:**

```
#include <stdio.h>
#include <stdlib.h>

wchar_t wchar = { 0x0073 };
char mbuffer[MB_CUR_MAX];

int main(void)
{
 int len;

 printf("Character encodings do %shave "
 "state-dependent \nencoding.\n",
 (wctomb(NULL, 0))
 ? "" : "not ");

 len = wctomb(mbuffer, wchar);
 mbuffer[len] = '\\0';
 printf("%s(%d)\n", mbuffer, len);
 return EXIT_SUCCESS;
}
```

This produces the output:

```
Character encodings do not have state-dependent
encoding.
s(1)
```

**Classification:**

ANSI

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

“String manipulation functions” and “Wide-character functions” in the summary of functions chapter.



## Synopsis:

```
#include <wctype.h>

wctrans_t wctrans(const char *property);
```

## Arguments:

*property*     The type of mapping; see below.

## Library:

`libc`

## Description:

The *wctrans()* function determines a mapping rule for wide-character codes according to the category `LC_CTYPE`, particularly for use with *towctrans()*.

The following mappings are defined in all locales, although additional classes may be defined for `LC_CTYPE`:

- `tolower`
- `toupper`

Use *setlocale()* to modify the category `LC_CTYPE`.

## Returns:

An object that you can use in a call to *towctrans()*, or 0 if the specified character mapping isn't valid for the current locale.

## Classification:

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*towctrans()*

String manipulation functions

Wide-character functions

## Synopsis:

```
#include <wctype.h>

wctype_t wctype(const char * property);
```

## Arguments:

*property*     A string that defines the property of the class; see below.

## Library:

`libc`

## Description:

The *wctype()* function determines a classification rule for wide-character codes according to the category `LC_CTYPE`, particularly for use with *iswctype()*.

Some classes are defined in all locales, although additional classes may be defined for `LC_CTYPE`. Use *setlocale()* to modify the category `LC_CTYPE`.

Defined Classes:

|       |       |        |
|-------|-------|--------|
| alnum | digit | punct  |
| alpha | graph | space  |
| blank | lower | upper  |
| cntrl | print | xdigit |

## Returns:

A `wctype_t` object that you can use in a call to *iswctype()*, or 0 if the character class name isn't valid for the current locale.

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*setlocale()*

“Character manipulation functions” and “Wide-character functions” in the summary of functions chapter.

## Synopsis:

```
#include <wchar.h>

wchar_t * wmemchr(const wchar_t * ws,
 wchar_t wc,
 size_t n);
```

## Arguments:

- ws*     The buffer that you want to search.
- wc*     The character that you're looking for.
- n*      The number of wide characters to search in the buffer.

## Library:

libc

## Description:

The *wmemchr()* function locates the first occurrence of *wc* in the first *n* wide characters of the buffer pointed to by *ws*.

The *wmemchr()* function is locale-independent and treats all **wchar\_t** values identically, even if they're null or invalid characters.

## Returns:

A pointer to the located character, or NULL if *wc* couldn't be found.

## Classification:

ANSI

### **Safety**

---

Cancellation point   No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*memccpy(), memcmp(), memcpy(), memicmp(), memmove(), memset()  
wcschr(), wcsrchr(), wmemcmp(), wmemcpy(), wmemmove(),  
wmemset()*

## Synopsis:

```
#include <wchar.h>

int wmemcmp(const wchar_t * ws1,
 const wchar_t * ws2,
 size_t n);
```

## Arguments:

*ws1*, *ws2*    The wide-character strings that you want to compare.

*n*            The number of wide characters to compare.

## Library:

libc

## Description:

The *wmemcmp()* function compares *n* wide characters of the buffer pointed to by *ws1* to those in the buffer pointed to by *ws2*.

## Returns:

<0    *ws1* is less than *ws2*.

0    *ws1* is equal to *ws2*.

>0    *ws1* is greater than *ws2*.

## Classification:

ANSI

### Safety

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*memcpy(), memcmp(), memchr(), memicmp(), memmove(), memset()  
wcsncmp(), wcsncpy(), wmemchr(), wmemcpy(), wmemmove(),  
wmemset()*



## Synopsis:

```
#include <wchar.h>

wchar_t * wmemcpy(wchar_t * ws1,
 const wchar_t * ws2,
 size_t n);
```

## Arguments:

- ws1* A pointer to the buffer that you want to copy the wide characters into.
- ws2* A pointer to the buffer that you want to copy the wide characters from.
- n* The number of wide characters to copy.

## Library:

`libc`

## Description:

The *wmemcpy()* function copies *n* wide characters from the buffer pointed to by *ws2* into the buffer pointed to by *ws1*.

The *wmemcpy()* function is locale-independent and treats all `wchar_t` values identically, even if they're null or invalid characters.



---

Copying overlapping buffers isn't guaranteed to work; use *wmemmove()* to copy buffers that overlap.

---

## Returns:

A pointer to the destination buffer (i.e the same pointer as *ws1*).

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memcpy()*, *memcmp()*, *memcpy()*, *memicmp()*, *memmove()*, *memset()*  
*wscpy()*, *wcsncpy()*, *wmemchr()*, *wmemcpy()*, *wmemmove()*,  
*wmemset()*

## Synopsis:

```
#include <wchar.h>

wchar_t * wmemmove(wchar_t * ws1,
 const wchar_t * ws2,
 size_t n);
```

## Arguments:

*ws1* A pointer to where you want the function to copy the data.

*ws2* A pointer to the buffer that you want to copy data from.

*n* The number of wide characters to copy.

## Library:

libc

## Description:

The *memmove()* function copies *n* wide characters from the buffer pointed to by *ws2* to the buffer pointed to by *ws1*. This function copies overlapping regions safely.

The *wmemmove()* function is locale-independent and treats all **wchar\_t** values identically, even if they're null or invalid characters.



---

Use *wmemcpy()* for greater speed when copying buffers that don't overlap.

---

## Returns:

A pointer to the destination buffer (i.e. the same pointed as *ws1*).

**Classification:**

ANSI

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | Yes |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*memcpy(), memcmp(), memcpy(), memicmp(), memmove(), memset()*  
*wmemchr(), wmemcmp(), wmemcpy(), wmemset()*

## Synopsis:

```
#include <wchar.h>

wchar_t * wmemset(wchar_t * ws,
 wchar_t wc,
 size_t n
```

## Arguments:

*ws*            A pointer to the memory that you want to set.

*wc*            The value that you want to store in each wide character.

*length*        The number of wide characters to set.

## Library:

`libc`

## Description:

The *memset()* function fills *n* wide characters starting at *ws* with the value *wc*.

The *wmemset()* function is locale-independent and treats all `wchar_t` values identically, even if they're null or invalid wide characters.

## Returns:

A pointer to the destination buffer (i.e. the same pointer as *ws*).

## Classification:

ANSI

### **Safety**

---

Cancellation point    No

*continued...*

**Safety**

---

|                   |     |
|-------------------|-----|
| Interrupt handler | Yes |
| Signal handler    | Yes |
| Thread            | Yes |

**See also:**

*memcpy(), memcmp(), memcpy(), memicmp(), memmove(), memset()  
wmemchr(), wmemcmp(), wmemcpy(), wmemmove()*

## Synopsis:

```
#include <wordexp.h>

int wordexp(const char * words,
 wordexp_t * pwordexp,
 int flags);
```

## Library:

libc

## Description:

The C bindings for performing word expansions aren't currently supported.

## Returns:

-1 to indicate an error (*errno* is set).

## Errors:

WRDE\_NOSYS

The *wordexp()* function isn't currently supported.

## Classification:

POSIX 1003.1a

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*glob(), globfree(), wordfree()*



## Synopsis:

```
#include <wordexp.h>

void wordfree(wordexp_t * pwordexp);
```

## Library:

libc

## Description:

The C bindings for performing word expansions aren't currently supported.

## Classification:

POSIX 1003.1a

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*glob(), globfree(), wordexp()*

# **wprintf()**

© 2004, QNX Software Systems Ltd.

*Write formatted output to stdout*

## **Synopsis:**

```
#include <wchar.h>

int wprintf(const char* format,
 ...);
```

## **Arguments:**

*format*     A wide-character string that specifies the format of the output. The formatting string determines what additional arguments you need to provide. For more information, see *printf()*.

## **Library:**

`libc`

## **Description:**

The *wprintf()* function writes output to the *stdout* stream, under control of the argument *format*. It's the wide-character version of *printf()*.

## **Returns:**

The number of characters written, or a negative value if an output error occurred (*errno* is set).

## **Classification:**

ANSI

### **Safety**

---

Cancellation point    Yes

Interrupt handler     No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*errno, fprintf(), fwprintf(), printf(), snprintf(), sprintf(), swprintf(),  
vfprintf(), vfwprintf(), vprintf(), vsnprintf(), vsprintf(), vswprintf(),  
vwprintf()*

## ***write()***

© 2004, QNX Software Systems Ltd.

*Write bytes to a file*

---

### **Synopsis:**

```
#include <unistd.h>

ssize_t write(int fd,
 const void* buf,
 size_t nbyte);
```

### **Arguments:**

*fd*      The file descriptor for the file you want to write in.

*buf*     A pointer to a buffer that contains the data you want to write.

*nbyte*   The number of bytes to write.

### **Library:**

`libc`

### **Description:**

The *write()* function attempts to write *nbyte* bytes to the file associated with the open file descriptor, *fd*, from the buffer pointed to by *buf*.

If *nbyte* is zero, *write()* returns zero, and has no other effect.

On a regular file or other file capable of seeking, and if `O_APPEND` isn't set, *write()* starts at a position in the file given by the file offset associated with *fd*. If `O_APPEND` is set, the file offset is set to the end of file before each write operation. Before successfully returning from *write()*, the file offset is incremented by the number of bytes actually written. On a regular file, if this incremented file offset is greater than the length of the file, the length of the file is set to this file offset.



---

Note that the *write()* call ignores advisory locks that may have been set by the *fcntl()* function.

---

On a file not capable of seeking, *write()* starts at the current position.

If *write()* requests that more bytes be written than there's room for (for example, all blocks on a disk are already allocated), only as many bytes as there's room for are written. For example, if there's only room for 80 more bytes in a file, a write of 512 bytes would return 80. The next write of a nonzero number of bytes would give a failure return (except as noted below).

When *write()* returns successfully, its return value is the number of bytes actually written to the file. This number is never greater than *nbyte*, although it may be less than *nbyte* under certain circumstances detailed below.

If *write()* is interrupted by a signal before it has written any data, it returns a value of -1, and *errno* is set to EINTR. However, if *write()* is interrupted by a signal after it has successfully written some data, it returns the number of bytes written.

If the value of *nbyte* is greater than INT\_MAX, *write()* returns -1 and sets *errno* to EINVAL. See <limits.h>.

Write requests to a pipe (or FIFO) are handled the same as a regular file, with the following exceptions:

- There's no file offset associated with a pipe, therefore each write request appends to the end of the pipe.
- Write requests of PIPE\_BUF bytes or less aren't interleaved with data from other processes doing writes on the same pipe. Writes of greater than PIPE\_BUF bytes may have data interleaved, on arbitrary boundaries, with writes by other processes, whether or not the O\_NONBLOCK flag is set.
- If the O\_NONBLOCK flag is clear, a write request may cause the process to block, but on normal completion it returns *nbyte*.

- If the `O_NONBLOCK` flag is set, write requests are handled differently, in the following ways:
  - The `write()` function doesn't block the process.
  - Write requests for `PIPE_BUF` bytes or less either succeed completely and return `nbyte`, or return -1 and `errno` is set to `EAGAIN`.

If you call `write()` with `nbyte` greater than `PIPE_BUF` bytes, it either transfers what it can and returns the number of bytes written, or transfers no data, returning -1 and setting `errno` to `EAGAIN`. Also, if `nbyte` is greater than `PIPE_BUF` bytes and all data previously written to the pipe has been read (that is, the pipe is empty), `write()` transfers at least `PIPE_BUF` bytes.

When attempting to write to a file (other than a pipe or FIFO) that supports nonblocking writes and can't accept the data immediately:

- If the `O_NONBLOCK` flag is clear, `write()` blocks until the data can be accepted.
- If the `O_NONBLOCK` flag is set, `write()` doesn't block the process. If some data can be written without blocking the process, `write()` transfers what it can and returns the number of bytes written. Otherwise, it returns -1 and sets `errno` to `EAGAIN`.

If `write()` is called with the file offset beyond the end-of-file, the file is extended to the current file offset with the intervening bytes filled with zeroes. This is a useful technique for pregrowing a file.

If `write()` succeeds, the `st_ctime` and `st_mtime` fields of the file are marked for update.

## Returns:

The number of bytes written, or -1 if an error occurred (`errno` is set).

## Errors:

|        |                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the write operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for writing.                                                                                       |
| EFBIG  | The file is a regular file, where <i>nbytes</i> is greater than 0, and the starting position is greater than or equal to the offset maximum associated with the file.    |
| EINTR  | The write operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EIO    | A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.                                                             |
| ENOSPC | There's no free space remaining on the device containing the file.                                                                                                       |
| ENOSYS | The <i>write()</i> function isn't implemented for the filesystem specified by <i>fdes</i> .                                                                              |
| EPIPE  | An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.                               |

## Examples:

```
#include <stdio.h>
#include <fcntl.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/stat.h>
#include <stdlib.h>

char buffer[] = { "A text record to be written" };

int main(void)
{
 int fd;
 int size_written;
```

```
/* open a file for output */
/* replace existing file if it exists */
fd = creat("myfile.dat", S_IRUSR | S_IWUSR);

/* write the text */
size_written = write(fd, buffer,
 sizeof(buffer));

/* test for error */
if(size_written != sizeof(buffer)) {
 perror("Error writing myfile.dat");
 return EXIT_FAILURE;
}

/* close the file */
close(fd);

return EXIT_SUCCESS;
}
```

**Classification:**

POSIX 1003.1

**Safety**

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close(), creat(), dup(), dup2(), errno, fcntl(), lseek(), open(), pipe(), read(), readv(), select(), writev()*



### **Synopsis:**

```
#include <unistd.h>

int writeblock(int fd,
 size_t blksize,
 unsigned block,
 int numblks,
 const void *buff);
```

### **Arguments:**

|                |                                                                                  |
|----------------|----------------------------------------------------------------------------------|
| <i>fd</i>      | The file descriptor for the file you want to write in.                           |
| <i>blksize</i> | The number of bytes in each block of data.                                       |
| <i>block</i>   | The block number from which to start writing. Blocks are numbered starting at 0. |
| <i>numblks</i> | The number of blocks to write.                                                   |
| <i>buff</i>    | A pointer to a buffer that contains the blocks of data that you want to write.   |

### **Library:**

`libc`

### **Description:**

The *writeblock()* function writes *numblks* blocks of data to the file associated with the open file descriptor, *fd*, from the buffer pointed to by *buff*, starting at block number *block*.

This function is useful for direct updating of raw blocks on a block special device (for example, raw disk blocks), but you can also use it for high-speed updating (for example, of database files). The speed gain is through the combined seek/write implicit in this call.

If *numblks* is zero, *writeblock()* returns zero, and has no other results.

If successful, *writeblock()* returns the number of blocks actually written to the disk associated with *fd*. This number is never greater than *numblks*, but could be less than *numblks* if one of the following occurs:

- The process attempts to write more blocks than implementation limits allow to be written in a single atomic operation.
- A write error occurred after writing at least one block, and you set one of the sync flags (O\_SYNC or O\_DSYNC — see *open()*) when you opened the file.

If a write error occurs on the first block and one of the sync flags is set, *writeblock()* returns -1 and sets *errno* to EIO.

If one of the sync flags is set, *writeblock()* doesn't return until the blocks are actually transferred to the disk. If neither of the flags is set, *writeblock()* places the blocks in the cache and schedules them for writing as soon as possible, but returns before the writing takes place.



---

In the latter instance, it's impossible for the application to know if the write succeeded or not (due to system failures or bad disk blocks). Using the sync flags significantly impacts the performance of *writeblock()*, but guarantees that the data can be recovered.

---

## Returns:

The number of blocks actually written. If an error occurred, *writeblock()* returns -1, sets *errno* to indicate the error, and doesn't change the contents of the buffer pointed to by *buff*.

## Errors:

- |       |                                                                                                       |
|-------|-------------------------------------------------------------------------------------------------------|
| EBADF | The <i>fd</i> argument isn't a valid file descriptor that's open for writing a block-oriented device. |
| EIO   | A physical write error occurred on the first block, and either O_DSYNC or O_SYNC is set.              |

EINVAL     The starting position is invalid (0 or negative), or beyond the end of the file.

## Classification:

QNX Neutrino

### Safety

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

## See also:

*open()*, *readblock()*, *write()*

## ***writev()***

© 2004, QNX Software Systems Ltd.

*Write bytes to a file*

### **Synopsis:**

```
#include <sys/uio.h>

ssize_t writev(int fildes,
 const iov_t* iov,
 int iovcnt);
```

### **Arguments:**

*fildes*     The file descriptor for the file you want to write in.

*iov*        An array of **iovt** objects that contain the data that you want to write.

*iovcnt*     The number of elements in the array.

### **Library:**

**libc**

### **Description:**

The *writev()* function performs the same action as *write()*, but gathers the output data from the *iovcnt* buffers specified by the members of the *iov* array: *iov[0]*, *iov[1]*, ..., *iov[iovcnt-1]*.

For *writev()*, the **iovt** structure contains the following members:

*iov\_base*     Base address of a memory area from which data should be written.

*iov\_len*       The length of the memory area.

The *writev()* function always writes a complete area before proceeding to the next.

The maximum number of entries in the *iov* array is **UIO\_MAXIOV**.



---

Note that *writev()* ignores advisory locks that may have been set by the *fcntl()* function.

---

If *writev()* is interrupted by a signal before it has written any data, it returns a value of -1, and *errno* is set to EINTR. However, if *writev()* is interrupted by a signal after it has successfully written some data, it will return the number of bytes written.

For more details, see the *write()* function.

## Returns:

The number of bytes written, or -1 if an error occurs (*errno* is set).

## Errors:

|        |                                                                                                                                                                          |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| EAGAIN | The O_NONBLOCK flag is set for the file descriptor, and the process would be delayed in the write operation.                                                             |
| EBADF  | The file descriptor, <i>fdes</i> , isn't a valid file descriptor open for writing.                                                                                       |
| EFBIG  | The file is a regular file, where <i>nbytes</i> is greater than 0, and the starting position is greater than or equal to the offset maximum associated with the file.    |
| EINTR  | The write operation was interrupted by a signal, and either no data was transferred, or the resource manager responsible for that file doesn't report partial transfers. |
| EINVAL | The <i>iovcnt</i> argument is less than or equal to 0, or greater than UIO_MAXIOV.                                                                                       |
| EIO    | A physical I/O error occurred (for example, a bad block on a disk). The precise meaning is device-dependent.                                                             |
| ENOSPC | There is no free space remaining on the device containing the file.                                                                                                      |
| ENOSYS | The <i>write()</i> function isn't implemented for the filesystem specified by <i>fdes</i> .                                                                              |

EPIPE      An attempt was made to write to a pipe (or FIFO) that isn't open for reading by any process. A SIGPIPE signal is also sent to the process.

**Classification:**

Standard Unix

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | Yes |
| Interrupt handler  | No  |
| Signal handler     | Yes |
| Thread             | Yes |

**See also:**

*close(), creat(), dup(), dup2(), errno, fcntl(), lseek(), open(), pipe(), read(), readv(), select(), write()*

**Synopsis:**

```
#include <wchar.h>

int wscanf(const char * format,
 ...);
```

**Arguments:**

*format*     A wide-character string that specifies the format of the input. For more information, see *scanf()*. The formatting string determines what additional arguments you need to provide.

**Library:**

libc

**Description:**

The *wscanf()* function scans input from *stdin* under control of the *format* argument, assigning values to the remaining arguments. It is the wide-character version of *scanf()* and uses the same conversions.

**Returns:**

The number of input arguments for which values were successfully scanned and stored, or EOF if the scanning reached the end of the input stream before storing any values.

**Classification:**

ANSI

**Safety**

---

Cancellation point    Yes

Interrupt handler      No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*fscanf()*, *fwscanf()*, *scanf()*, *sscanf()*, *swscanf()*, *vfscanf()*, *vwscanf()*,  
*vscanf()*, *vsscanf()*, *vswscanf()*, *vwscanf()*



## Synopsis:

```
#include <math.h>

double y0(double x);

float y0f(float x);
```

## Arguments:

*x* The number that you want to compute the Bessel function for.

## Library:

`libbessel`

## Description:

Compute the Bessel function of the second kind for *x*.

## Returns:

The result of the Bessel function of *x*.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Classification:

*y0()* is standard Unix; *y0f()* is ANSI (draft)

### Safety

---

Cancellation point No

Interrupt handler No

*continued...*

**Safety**

---

|                |     |
|----------------|-----|
| Signal handler | No  |
| Thread         | Yes |

**See also:**

*errno, j0(), j1(), jn(), y1(), yn()*

## Synopsis:

```
#include <math.h>

double y1(double x);

float y1f(float x);
```

## Arguments:

*x* The number that you want to compute the Bessel function for.

## Library:

libbessel

## Description:

Compute the Bessel function of the second kind for *x*.

## Returns:

The result of the Bessel function of *x*.



---

If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Examples:

```
#include <stdlib.h>
#include <stdio.h>
#include <math.h>

int main(void)
{
 double x, y, z;

 x = j0(2.4);
 y = y1(1.58);
 z = jn(3, 2.4);
```

```
printf("j0(2.4) = %f, y1(1.58) = %f\n", x, y);
printf("jn(3,2.4) = %f\n", z);

return EXIT_SUCCESS;
}
```

**Classification:**

*y1()* is standard Unix; *y1f()* is ANSI (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno, j0(), j1(), jn(), y0(), yn()*

## Synopsis:

```
#include <math.h>

double yn(int n,
 double x);

float ynf(int n,
 float x);
```

## Arguments:

*n, x*    The numbers that you want to compute the Bessel function for.

## Library:

libbessel

## Description:

Compute the Bessel function of the second kind for *n* and *x*.

## Returns:

The result of the Bessel function of *n* and *x*.



If an error occurs, these functions return 0, but this is also a valid mathematical result. If you want to check for errors, set *errno* to 0, call the function, and then check *errno* again. These functions don't change *errno* if no errors occurred.

---

## Classification:

*yn()* is standard Unix; *ynf()* is ANSI (draft)

**Safety**

---

|                    |     |
|--------------------|-----|
| Cancellation point | No  |
| Interrupt handler  | No  |
| Signal handler     | No  |
| Thread             | Yes |

**See also:**

*errno*, *j0()*, *j1()*, *jn()*, *y0()*, *y1()*

## *Appendix A*

---

# **SOCKS — A Basic Firewall**

### *In this appendix...*

About SOCKS 3571  
How to SOCKSify a client 3571  
What SOCKS expects 3572





## About SOCKS

SOCKS is a package consisting of a proxy server, client programs (`rftp` and `rtelnet`), and a library (`libsocks`) for adapting other applications into new client programs.

The original SOCKS was written by David Koblas (`koblas@netcom.com`). The SOCKS protocol has changed over time. The client library shipped as of printing corresponds to SOCKS v4.2. Since the server and the clients must use the same SOCKS protocol, this library doesn't work with servers of previous releases; clients compiled with these libraries won't work with older servers.

## How to SOCKSify a client



If your client is using UDP to transfer data, you can't use SOCKS. To see if your client uses UDP, search for the string "SOCK\_DGRAM" in your source.

- 1 At or near the beginning of `main()`, you can add a call to `SOCKSinit()`.

You can omit this step; the only reason for calling `SOCKSinit()` directly is to associate a name with your SOCKS client (rather than the generic "SOCKSclient" default string).

- 2 Add the following options to your compile commands:
 

```
-Dconnect=Rconnect -Dgetsockname=Rgetsockname \
-Dbind=Rbind -Daccept=Raccept -Dlisten=Rlisten \
-Drcmd=Rrcmd -Dselect=Rselect
```

If you're using a **Makefile**, add these options to the definition of macro `CFLAGS`.

These options replace calls to certain functions with versions that use the SOCKS server:

| <b>Non-SOCKS function:</b> | <b>SOCKS function:</b> |
|----------------------------|------------------------|
| <i>accept()</i>            | <i>Raccept()</i>       |
| <i>bind()</i>              | <i>Rbind()</i>         |
| <i>connect()</i>           | <i>Rconnect()</i>      |
| <i>getsockname()</i>       | <i>Rgetsockname()</i>  |
| <i>listen()</i>            | <i>Rlisten()</i>       |
| <i>rcmd()</i>              | <i>Rrcmd()</i>         |
| <i>select()</i>            | <i>Rselect()</i>       |

- 3 Link against the SOCKS library by adding `-l socks` to your link line.

If you're using a **Makefile**, simply add this information to the definition of the macro `LDFLAGS`.

For most programs, the above steps should be sufficient to SOCKSify the package. If the above doesn't work, you may need to look at things a little more closely. The next section describes how the SOCKS library expects to be used.

## What SOCKS expects

The SOCKS library covers only some of the socket functions, which must be called in a particular order:




---

You must use TCP; SOCKS doesn't support UDP.

---

- 1 The first socket function invoked must be either *connect()* or *rcmd()*.
- 2 If you call *connect()* on a nonblocking socket, no I/O can occur on that socket until another *connect()*, with the same arguments, returns -1 and sets *errno* to `EISCONN`. This is required even if you use *select()* on write to check the readiness of that socket.



---

While a connection is still pending, don't try to start another connection via *connect()*, or start a sequence of *bind()*, *getsockname()*, *listen()*, and *accept()*.

---

- 3 You must call *bind()* after a successful *connect()* call to a host for a specific service.
- 4 You must follow the call to *bind()* by calls to *getsockname()*, *listen()*, and *accept()*, in that order.

Most client programs fit these assumptions very well and can be SOCKSified without changing the code at all using the steps described in “How to SOCKSify a client.”

Some client programs use a *bind()* before each *connect()*. If the *bind()* is used to claim a specific port or a specific network interface, the current SOCKS library can't accommodate such use. Very often though, such a *bind()* call is there for no specific reason and may simply be deleted.



***Appendix B***

---

**Third-Party Copyright Notices**



## **BSD Stack**

**Copyright © 1997 Christopher G. Demetriou.**

**All rights reserved.**

**Copyright © 1982, 1986, 1989, 1991, 1993 The Regents of the University of California.**

**All rights reserved.**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1** Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2** Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3** Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT

OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## **BSD Stack and Various Utilities**

**Copyright © 1998 The NetBSD Foundation, Inc.**

**All rights reserved.**

**This code is derived from software contributed to The NetBSD Foundation by Public Access Networks Corporation ("Panix"). It was developed under contract to Panix by Eric Haszlkiewicz and Thor Lancelot Simon.**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
- 4 Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE



FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Copyright © 1995 The NetBSD Foundation, Inc. All rights reserved.**

This code is derived from software contributed to The NetBSD Foundation by Christos Zoulas. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
- 4 Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT

LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Copyright © 1996, 1997 The NetBSD Foundation, Inc. All rights reserved.**

This code is derived from software contributed to The NetBSD Foundation by Jason R. Thorpe of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by the NetBSD Foundation, Inc. and its contributors.
- 4 Neither the name of The NetBSD Foundation nor the names of its contributors may be used to endorse or promote products

derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Copyright © 1996 Matt Thomas matt@3am-software.com.  
All rights reserved.**

**Copyright © 1982, 1986, 1988, 1993 The Regents of the University of California.  
All rights reserved.**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 Neither the name of the University nor the names of its contributors may be used to endorse or promote products

derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Portions Copyright © 1993 by Digital Equipment Corporation.**

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of Digital Equipment Corporation not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED “AS IS” AND DIGITAL EQUIPMENT CORP. DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL DIGITAL EQUIPMENT CORPORATION BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT

OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

**Portions Copyright © 1995 by International Business Machines, Inc.**

International Business Machines, Inc. (hereinafter called IBM) grants permission under its copyrights to use, copy, modify, and distribute this Software with or without fee, provided that the above copyright notice and all paragraphs of this notice appear in all copies, and that the name of IBM not be used in connection with the marketing of any product incorporating the Software or modifications thereof, without specific, written prior permission.

To the extent it has a right to do so, IBM grants an immunity from suit under its patents, if any, for the use, sale or manufacture of products to the extent that such products are used for performing Domain Name System dynamic updates in TCP/IP networks by means of the Software. No immunity is granted for any product per se or for any other function of any product.

THE SOFTWARE IS PROVIDED "AS IS", AND IBM DISCLAIMS ALL WARRANTIES, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT SHALL IBM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE, EVEN IF IBM IS APPRISED OF THE POSSIBILITY OF SUCH DAMAGES.

**Copyright © 1996 by Internet Software Consortium.**

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND

FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

**All of the documentation and software included in the third BSD Networking Software Release is copyrighted by The Regents of the University of California.**

**Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1993  
The Regents of the University of California.  
All rights reserved.**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE

LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## **MINIX Operating System**

**Copyright © 1987,1997**

**Prentice Hall**

**All rights reserved.**

Redistribution and use of the MINIX operating system in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 Neither the name of Prentice Hall nor the names of the software authors or contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS, AUTHORS, AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL

PRENTICE HALL OR ANY AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Regular Expression Handling

**Copyright © 1992, 1993, 1994, 1997**

**Henry Spencer.**

**All rights reserved.**

This software is not subject to any license of the American Telephone and Telegraph Company or of the Regents of the University of California.

Permission is granted to anyone to use this software for any purpose on any computer system, and to alter it and redistribute it, subject to the following restrictions:

- 1 The author is not responsible for the consequences of use of this software, no matter how awful, even if they arise from flaws in it.
- 2 The origin of this software must not be misrepresented, either by explicit claim or by omission. Since few users ever read sources, credits must appear in the documentation.
- 3 Altered versions must be plainly marked as such, and must not be misrepresented as being the original software. Since few users ever read sources, credits must appear in the documentation.
- 4 This notice may not be removed or altered.



## **Remote Procedure Call (RPC)**

**Copyright © 1984, 1985, 1986, 1987**  
**Sun Microsystems, Inc.**  
**2550 Garcia Avenue**  
**Mountain View, California 94043**

Sun RPC is a product of Sun Microsystems, Inc. and is provided for unrestricted use provided that this legend is included on all tape media and as a part of the software program in whole or part. Users may copy or modify Sun RPC without charge, but are not authorized to license or distribute it to anyone else except as part of a product or program developed by the user.

SUN RPC IS PROVIDED AS IS WITH NO WARRANTIES OF ANY KIND INCLUDING THE WARRANTIES OF DESIGN, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE.

Sun RPC is provided with no support and without any obligation on the part of Sun Microsystems, Inc. to assist in its use, correction, modification or enhancement.

SUN MICROSYSTEMS, INC. SHALL HAVE NO LIABILITY WITH RESPECT TO THE INFRINGEMENT OF COPYRIGHTS, TRADE SECRETS OR ANY PATENTS BY SUN RPC OR ANY PART THEREOF.

In no event will Sun Microsystems, Inc. be liable for any lost revenue or profits or other special, indirect and consequential damages, even if Sun has been advised of the possibility of such damages.

## **SNMPv2**

**Copyright © 1988, 1989, 1991**  
**Carnegie Mellon University**  
**All Rights Reserved**

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted,

provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Carnegie Mellon University not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

## SOCKS

**Copyright © 1989**  
**The Regents of the University of California.**  
**All rights reserved.**

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- 1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- 2 Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- 3 Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

**Portions Copyright © 1993, 1994  
by NEC Systems Laboratory.**

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies, and that the name of NEC Systems Laboratory not be used in advertising or publicity pertaining to distribution of the document or software without specific, written prior permission.

THE SOFTWARE IS PROVIDED "AS IS" AND NEC SYSTEMS LABORATORY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NEC SYSTEMS LABORATORY BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.



## ***Appendix C***

---

### **Summary of Safety Information**

#### ***In this appendix...***

Cancellation points 3593  
Interrupt handlers 3598  
Signal handlers 3601  
Multithreaded programs 3614



## Cancellation points

The following functions are cancellation points:

|                              |                           |
|------------------------------|---------------------------|
| <i>ConnectAttach()</i>       | <i>closelog()</i>         |
| <i>ConnectAttach_r()</i>     | <i>connect()</i>          |
| <i>ConnectDetach()</i>       | <i>creat()</i>            |
| <i>ConnectDetach_r()</i>     | <i>creat64()</i>          |
| <i>ConnectServerInfo()</i>   | <i>delay()</i>            |
| <i>ConnectServerInfo_r()</i> | <i>devctl()</i>           |
| <i>InterruptWait()</i>       | <i>dispatch_block()</i>   |
| <i>InterruptWait_r()</i>     | <i>dispatch_unblock()</i> |
| <i>MsgSend()</i>             | <i>dlclose()</i>          |
| <i>MsgSend_r()</i>           | <i>dlopen()</i>           |
| <i>MsgSendsv()</i>           | <i>ds_clear()</i>         |
| <i>MsgSendsv_r()</i>         | <i>ds_create()</i>        |
| <i>MsgSendv()</i>            | <i>ds_deregister()</i>    |
| <i>MsgSendv_r()</i>          | <i>ds_flags()</i>         |
| <i>MsgSendvs()</i>           | <i>ds_get()</i>           |
| <i>MsgSendvs_r()</i>         | <i>ds_register()</i>      |
| <i>SignalSuspend()</i>       | <i>ds_set()</i>           |
| <i>SignalSuspend_r()</i>     | <i>endgrent()</i>         |
| <i>SignalWaitinfo()</i>      | <i>endhostent()</i>       |
| <i>SignalWaitinfo_r()</i>    | <i>endnetent()</i>        |
| <i>SyncCondvarSignal()</i>   | <i>endprotoent()</i>      |
| <i>SyncCondvarSignal_r()</i> | <i>endpwent()</i>         |
| <i>SyncCondvarWait()</i>     | <i>endservent()</i>       |
| <i>SyncCondvarWait_r()</i>   | <i>endspent()</i>         |
| <i>ThreadJoin()</i>          | <i>endutent()</i>         |
| <i>ThreadJoin_r()</i>        | <i>eof()</i>              |
| <i>accept()</i>              | <i>err()</i>              |
| <i>aio_suspend()</i>         | <i>errx()</i>             |
| <i>cfgopen()</i>             | <i>fcfgopen()</i>         |
| <i>chsize()</i>              | <i>fchown()</i>           |
| <i>clock_nanosleep()</i>     | <i>fclose()</i>           |
| <i>close()</i>               | <i>fcloseall()</i>        |
| <i>closedir()</i>            | <i>fdopen()</i>           |

|                           |                          |
|---------------------------|--------------------------|
| <i>fflush()</i>           | <i>getgrgid()</i>        |
| <i>fgetc()</i>            | <i>getgrgid_r()</i>      |
| <i>fgetchar()</i>         | <i>getgrnam()</i>        |
| <i>fgets()</i>            | <i>getgrnam_r()</i>      |
| <i>fgetspent()</i>        | <i>getgrouplist()</i>    |
| <i>fgetwc()</i>           | <i>gethostbyaddr()</i>   |
| <i>fgetwts()</i>          | <i>gethostbyaddr_r()</i> |
| <i>flushall()</i>         | <i>gethostbyname()</i>   |
| <i>fopen()</i>            | <i>gethostbyname2()</i>  |
| <i>fopen64()</i>          | <i>gethostbyname_r()</i> |
| <i>forkpty()</i>          | <i>gethostent()</i>      |
| <i>fprintf()</i>          | <i>gethostent_r()</i>    |
| <i>fputc()</i>            | <i>gethostname()</i>     |
| <i>fputchar()</i>         | <i>getifaddrs()</i>      |
| <i>fputs()</i>            | <i>getlogin()</i>        |
| <i>fputwc()</i>           | <i>getlogin_r()</i>      |
| <i>fputwts()</i>          | <i>getnameinfo()</i>     |
| <i>fread()</i>            | <i>getnetbyaddr()</i>    |
| <i>freopen()</i>          | <i>getnetbyname()</i>    |
| <i>freopen64()</i>        | <i>getnetent()</i>       |
| <i>fscanf()</i>           | <i>getopt()</i>          |
| <i>fsync()</i>            | <i>getpass()</i>         |
| <i>ftell()</i>            | <i>getpeername()</i>     |
| <i>ftello()</i>           | <i>getprotobyname()</i>  |
| <i>ftw()</i>              | <i>getprotobyname()</i>  |
| <i>ftw64()</i>            | <i>getprotoent()</i>     |
| <i>fwide()</i>            | <i>getpwent()</i>        |
| <i>fwprintf()</i>         | <i>getpwnam()</i>        |
| <i>fwrite()</i>           | <i>getpwnam_r()</i>      |
| <i>fwscanf()</i>          | <i>getpwuid()</i>        |
| <i>getaddrinfo()</i>      | <i>getpwuid_r()</i>      |
| <i>getc()</i>             | <i>gets()</i>            |
| <i>getc_unlocked()</i>    | <i>getservbyname()</i>   |
| <i>getchar()</i>          | <i>getservbyport()</i>   |
| <i>getchar_unlocked()</i> | <i>getservent()</i>      |
| <i>getcwd()</i>           | <i>getsockname()</i>     |
| <i>getgrent()</i>         | <i>getsockopt()</i>      |



|                              |                            |
|------------------------------|----------------------------|
| <i>getspent()</i>            | <i>mq_timedreceive()</i>   |
| <i>getspent_r()</i>          | <i>mq_timedsend()</i>      |
| <i>getspnam()</i>            | <i>msync()</i>             |
| <i>getspnam_r()</i>          | <i>name_attach()</i>       |
| <i>getutent()</i>            | <i>name_close()</i>        |
| <i>getutid()</i>             | <i>name_detach()</i>       |
| <i>getutline()</i>           | <i>name_open()</i>         |
| <i>getw()</i>                | <i>nanosleep()</i>         |
| <i>getwc()</i>               | <i>nap()</i>               |
| <i>getwchar()</i>            | <i>napms()</i>             |
| <i>getwd()</i>               | <i>nbaconnect()</i>        |
| <i>glob()</i>                | <i>nbaconnect_result()</i> |
| <i>herror()</i>              | <i>netmgr_ndtostr()</i>    |
| <i>if_indebyname()</i>       | <i>netmgr_strtond()</i>    |
| <i>if_nameindex()</i>        | <i>nftw()</i>              |
| <i>if_nametoindex()</i>      | <i>nftw64()</i>            |
| <i>initstate()</i>           | <i>open()</i>              |
| <i>input_line()</i>          | <i>open64()</i>            |
| <i>iofunc_attr_lock()</i>    | <i>opendir()</i>           |
| <i>iofunc_attr_trylock()</i> | <i>openfd()</i>            |
| <i>isfdtype()</i>            | <i>openlog()</i>           |
| <i>listen()</i>              | <i>openpty()</i>           |
| <i>login_tty()</i>           | <i>pathfind()</i>          |
| <i>ltrunc()</i>              | <i>pathfind_r()</i>        |
| <i>message_attach()</i>      | <i>pathmgr_symlink()</i>   |
| <i>message_connect()</i>     | <i>pathmgr_unlink()</i>    |
| <i>message_detach()</i>      | <i>pause()</i>             |
| <i>mknod()</i>               | <i>pccard_arm()</i>        |
| <i>mkstemp()</i>             | <i>pccard_attach()</i>     |
| <i>mktemp()</i>              | <i>pccard_detach()</i>     |
| <i>modem_open()</i>          | <i>pccard_info()</i>       |
| <i>modem_read()</i>          | <i>pccard_lock()</i>       |
| <i>modem_script()</i>        | <i>pccard_raw_read()</i>   |
| <i>modem_write()</i>         | <i>pccard_unlock()</i>     |
| <i>mount()</i>               | <i>pci_attach()</i>        |
| <i>mq_receive()</i>          | <i>pci_attach_device()</i> |
| <i>mq_send()</i>             | <i>pci_detach()</i>        |

|                                     |                                |
|-------------------------------------|--------------------------------|
| <i>pci_detach_device()</i>          | <i>putc_unlocked()</i>         |
| <i>pci_find_class()</i>             | <i>putchar()</i>               |
| <i>pci_find_device()</i>            | <i>putchar_unlocked()</i>      |
| <i>pci_irq_routing_options()</i>    | <i>puts()</i>                  |
| <i>pci_map_irq()</i>                | <i>putspent()</i>              |
| <i>pci_present()</i>                | <i>pututline()</i>             |
| <i>pci_read_config()</i>            | <i>putw()</i>                  |
| <i>pci_read_config16()</i>          | <i>putwc()</i>                 |
| <i>pci_read_config32()</i>          | <i>putwchar()</i>              |
| <i>pci_read_config8()</i>           | <i>pwrite()</i>                |
| <i>pci_rescan_bus()</i>             | <i>pwrite64()</i>              |
| <i>pci_write_config()</i>           | <i>rcmd()</i>                  |
| <i>pci_write_config16()</i>         | <i>read()</i>                  |
| <i>pci_write_config32()</i>         | <i>read_main_config_file()</i> |
| <i>pci_write_config8()</i>          | <i>readblock()</i>             |
| <i>pclose()</i>                     | <i>readcond()</i>              |
| <i>perror()</i>                     | <i>readdir()</i>               |
| <i>poll()</i>                       | <i>readv()</i>                 |
| <i>popen()</i>                      | <i>realpath()</i>              |
| <i>pread()</i>                      | <i>recv()</i>                  |
| <i>pread64()</i>                    | <i>recvfrom()</i>              |
| <i>printf()</i>                     | <i>recvmsg()</i>               |
| <i>pthread_cond_timedwait()</i>     | <i>remove()</i>                |
| <i>pthread_cond_wait()</i>          | <i>rename()</i>                |
| <i>pthread_join()</i>               | <i>res_init()</i>              |
| <i>pthread_rwlock_rdlock()</i>      | <i>res_mkquery()</i>           |
| <i>pthread_rwlock_timedrdlock()</i> | <i>res_query()</i>             |
| <i>pthread_rwlock_timedwrlock()</i> | <i>res_querydomain()</i>       |
| <i>pthread_rwlock_tryrdlock()</i>   | <i>res_search()</i>            |
| <i>pthread_rwlock_trywrlock()</i>   | <i>res_send()</i>              |
| <i>pthread_rwlock_wrlock()</i>      | <i>resmgr_attach()</i>         |
| <i>pthread_sleepon_lock()</i>       | <i>resmgr_block()</i>          |
| <i>pthread_sleepon_timedwait()</i>  | <i>resmgr_context_alloc()</i>  |
| <i>pthread_sleepon_wait()</i>       | <i>resmgr_context_free()</i>   |
| <i>pthread_testcancel()</i>         | <i>resmgr_detach()</i>         |
| <i>pulse_attach()</i>               | <i>resmgr_devino()</i>         |
| <i>putc()</i>                       | <i>resmgr_handler()</i>        |

|                                 |                              |
|---------------------------------|------------------------------|
| <i>rewind()</i>                 | <i>shm_ctl()</i>             |
| <i>rewinddir()</i>              | <i>shutdown()</i>            |
| <i>rresvport()</i>              | <i>sigpause()</i>            |
| <i>rsrddbmgr_attach()</i>       | <i>sigsuspend()</i>          |
| <i>rsrddbmgr_create()</i>       | <i>sigtimedwait()</i>        |
| <i>rsrddbmgr_destroy()</i>      | <i>sigwait()</i>             |
| <i>rsrddbmgr_detach()</i>       | <i>sigwaitinfo()</i>         |
| <i>rsrddbmgr_devno_attach()</i> | <i>sleep()</i>               |
| <i>rsrddbmgr_devno_detach()</i> | <i>slogb()</i>               |
| <i>rsrddbmgr_query()</i>        | <i>slogf()</i>               |
| <i>ruserok()</i>                | <i>slogi()</i>               |
| <i>scandir()</i>                | <i>snmp_close()</i>          |
| <i>scanf()</i>                  | <i>snmp_open()</i>           |
| <i>sctp_bindx()</i>             | <i>snmp_read()</i>           |
| <i>sctp_connectx()</i>          | <i>snmp_send()</i>           |
| <i>sctp_getladdrs()</i>         | <i>snmp_timeout()</i>        |
| <i>sctp_getpaddrs()</i>         | <i>socketmark()</i>          |
| <i>sctp_peeloff()</i>           | <i>socket()</i>              |
| <i>sctp_rcvmsg()</i>            | <i>socketpair()</i>          |
| <i>sctp_sendmsg()</i>           | <i>sopen()</i>               |
| <i>seekdir()</i>                | <i>sopenfd()</i>             |
| <i>select_attach()</i>          | <i>sysctl()</i>              |
| <i>select_detach()</i>          | <i>syslog()</i>              |
| <i>sem_timedwait()</i>          | <i>system()</i>              |
| <i>sem_wait()</i>               | <i>tcdrain()</i>             |
| <i>send()</i>                   | <i>tell()</i>                |
| <i>sendmsg()</i>                | <i>tell64()</i>              |
| <i>sendto()</i>                 | <i>telldir()</i>             |
| <i>setgrent()</i>               | <i>thread_pool_control()</i> |
| <i>setgroups()</i>              | <i>thread_pool_destroy()</i> |
| <i>sethostname()</i>            | <i>thread_pool_limits()</i>  |
| <i>setnetent()</i>              | <i>thread_pool_start()</i>   |
| <i>setprotoent()</i>            | <i>tmpfile()</i>             |
| <i>setpwent()</i>               | <i>tmpfile64()</i>           |
| <i>setservent()</i>             | <i>tmpnam()</i>              |
| <i>setsockopt()</i>             | <i>truncate()</i>            |
| <i>setutent()</i>               | <i>umount()</i>              |

|                   |                     |
|-------------------|---------------------|
| <i>ungetc()</i>   | <i>vwscanf()</i>    |
| <i>ungetwc()</i>  | <i>wait()</i>       |
| <i>unlink()</i>   | <i>wait3()</i>      |
| <i>usleep()</i>   | <i>wait4()</i>      |
| <i>utmpname()</i> | <i>waitid()</i>     |
| <i>verr()</i>     | <i>waitpid()</i>    |
| <i>verrx()</i>    | <i>warn()</i>       |
| <i>vfscanf()</i>  | <i>warnx()</i>      |
| <i>vwscanf()</i>  | <i>wordexp()</i>    |
| <i>vscanf()</i>   | <i>wprintf()</i>    |
| <i>vslogf()</i>   | <i>write()</i>      |
| <i>vsyslog()</i>  | <i>writeblock()</i> |
| <i>vwarn()</i>    | <i>writev()</i>     |
| <i>vwarnx()</i>   | <i>wscanf()</i>     |

See the “Caveats” section for the following functions for more information:

|                           |                   |
|---------------------------|-------------------|
| <i>dispatch_handler()</i> | <i>spawnlpe()</i> |
| <i>fcntl()</i>            | <i>spawnv()</i>   |
| <i>spawnl()</i>           | <i>spawnve()</i>  |
| <i>spawnle()</i>          | <i>spawnvp()</i>  |
| <i>spawnlp()</i>          | <i>spawnvpe()</i> |

## Interrupt handlers

It’s safe to call the following functions from an interrupt handler:

|                       |                           |
|-----------------------|---------------------------|
| <i>ENDIAN_BE16()</i>  | <i>ENDIAN_SWAP16()</i>    |
| <i>ENDIAN_BE32()</i>  | <i>ENDIAN_SWAP32()</i>    |
| <i>ENDIAN_BE64()</i>  | <i>ENDIAN_SWAP64()</i>    |
| <i>ENDIAN_LE16()</i>  | <i>GETIOVBASE()</i>       |
| <i>ENDIAN_LE32()</i>  | <i>GETIOVLEN()</i>        |
| <i>ENDIAN_LE64()</i>  | <i>InterruptDisable()</i> |
| <i>ENDIAN_RET16()</i> | <i>InterruptEnable()</i>  |
| <i>ENDIAN_RET32()</i> | <i>InterruptLock()</i>    |
| <i>ENDIAN_RET64()</i> | <i>InterruptMask()</i>    |

|                              |                              |
|------------------------------|------------------------------|
| <i>InterruptUnlock()</i>     | <i>gai_strerror()</i>        |
| <i>InterruptUnmask()</i>     | <i>htonl()</i>               |
| <i>ND_NODE_CMP()</i>         | <i>htons()</i>               |
| <i>SETIOV()</i>              | <i>hwi_find_item()</i>       |
| <i>SYSPAGE_CPU_ENTRY()</i>   | <i>hwi_find_tag()</i>        |
| <i>SYSPAGE_ENTRY()</i>       | <i>hwi_off2tag()</i>         |
| <i>UNALIGNED_PUT16()</i>     | <i>hwi_tag2off()</i>         |
| <i>UNALIGNED_PUT32()</i>     | <i>in16()</i>                |
| <i>UNALIGNED_PUT64()</i>     | <i>in16s()</i>               |
| <i>UNALIGNED_RET16()</i>     | <i>in32()</i>                |
| <i>UNALIGNED_RET32()</i>     | <i>in32s()</i>               |
| <i>UNALIGNED_RET64()</i>     | <i>in8()</i>                 |
| <i>_RESMGR_NPARTS()</i>      | <i>in8s()</i>                |
| <i>_RESMGR_PTR()</i>         | <i>inbe16()</i>              |
| <i>_RESMGR_STATUS()</i>      | <i>inbe32()</i>              |
| <i>abs()</i>                 | <i>inle16()</i>              |
| <i>alphasort()</i>           | <i>inle32()</i>              |
| <i>atoh()</i>                | <i>ipsec_get_policylen()</i> |
| <i>atoi()</i>                | <i>ipsec_strerror()</i>      |
| <i>atol()</i>                | <i>isalnum()</i>             |
| <i>atoll()</i>               | <i>isalpha()</i>             |
| <i>atomic_add()</i>          | <i>isascii()</i>             |
| <i>atomic_add_value()</i>    | <i>iscntrl()</i>             |
| <i>atomic_clr()</i>          | <i>isdigit()</i>             |
| <i>atomic_clr_value()</i>    | <i>isgraph()</i>             |
| <i>atomic_set()</i>          | <i>islower()</i>             |
| <i>atomic_set_value()</i>    | <i>isprint()</i>             |
| <i>atomic_sub()</i>          | <i>ispunct()</i>             |
| <i>atomic_sub_value()</i>    | <i>isspace()</i>             |
| <i>atomic_toggle()</i>       | <i>isupper()</i>             |
| <i>atomic_toggle_value()</i> | <i>iswalnum()</i>            |
| <i>basename()</i>            | <i>iswalpha()</i>            |
| <i>bcmp()</i>                | <i>iswcntrl()</i>            |
| <i>bcopy()</i>               | <i>iswdigit()</i>            |
| <i>bsearch()</i>             | <i>iswgraph()</i>            |
| <i>bzero()</i>               | <i>iswlower()</i>            |
| <i>div()</i>                 | <i>iswprint()</i>            |

|                         |                      |
|-------------------------|----------------------|
| <i>iswpunct()</i>       | <i>sigdelset()</i>   |
| <i>iswspace()</i>       | <i>sigemptyset()</i> |
| <i>iswupper()</i>       | <i>sigfillset()</i>  |
| <i>iswxdigit()</i>      | <i>sigismember()</i> |
| <i>isxdigit()</i>       | <i>sigmask()</i>     |
| <i>itoa()</i>           | <i>straddstr()</i>   |
| <i>lltoa()</i>          | <i>strcasecmp()</i>  |
| <i>lsearch()</i>        | <i>strcat()</i>      |
| <i>ltoa()</i>           | <i>strchr()</i>      |
| <i>max()</i>            | <i>strcmp()</i>      |
| <i>memccpy()</i>        | <i>strcmpi()</i>     |
| <i>memchr()</i>         | <i>strcoll()</i>     |
| <i>memcmp()</i>         | <i>strcpy()</i>      |
| <i>memcpy()</i>         | <i>strcspn()</i>     |
| <i>memcpyv()</i>        | <i>stricmp()</i>     |
| <i>memicmp()</i>        | <i>strlen()</i>      |
| <i>memmove()</i>        | <i>strlwr()</i>      |
| <i>memset()</i>         | <i>strncasecmp()</i> |
| <i>min()</i>            | <i>strncat()</i>     |
| <i>nanospin_count()</i> | <i>strncmp()</i>     |
| <i>nsec2timespec()</i>  | <i>strncpy()</i>     |
| <i>ntohl()</i>          | <i>strnicmp()</i>    |
| <i>ntohs()</i>          | <i>strnset()</i>     |
| <i>offsetof()</i>       | <i>strpbrk()</i>     |
| <i>out16()</i>          | <i>strrchr()</i>     |
| <i>out16s()</i>         | <i>strrev()</i>      |
| <i>out32()</i>          | <i>strsep()</i>      |
| <i>out32s()</i>         | <i>strset()</i>      |
| <i>out8()</i>           | <i>strspn()</i>      |
| <i>out8s()</i>          | <i>strstr()</i>      |
| <i>outbe16()</i>        | <i>strtoimax()</i>   |
| <i>outbe32()</i>        | <i>strtok_r()</i>    |
| <i>outle16()</i>        | <i>strtol()</i>      |
| <i>outle32()</i>        | <i>strtoll()</i>     |
| <i>rindex()</i>         | <i>strtoul()</i>     |
| <i>setdomainname()</i>  | <i>strtoull()</i>    |
| <i>sigaddset()</i>      | <i>strtoumax()</i>   |

|                        |                    |
|------------------------|--------------------|
| <i>strupr()</i>        | <i>wcslen()</i>    |
| <i>strxfrm()</i>       | <i>wcsncat()</i>   |
| <i>swab()</i>          | <i>wcsncmp()</i>   |
| <i>timespec2nsec()</i> | <i>wcsncpy()</i>   |
| <i>tolower()</i>       | <i>wcspbrk()</i>   |
| <i>toupper()</i>       | <i>wcsrchr()</i>   |
| <i>towctrans()</i>     | <i>wcsspn()</i>    |
| <i>towlower()</i>      | <i>wcsstr()</i>    |
| <i>towupper()</i>      | <i>wcstoimax()</i> |
| <i>ulltoa()</i>        | <i>wcstol()</i>    |
| <i>ultoa()</i>         | <i>wcstoll()</i>   |
| <i>utoa()</i>          | <i>wcstoul()</i>   |
| <i>va_arg()</i>        | <i>wcstoull()</i>  |
| <i>va_copy()</i>       | <i>wcstoumax()</i> |
| <i>va_end()</i>        | <i>wctrans()</i>   |
| <i>va_start()</i>      | <i>wctype()</i>    |
| <i>wscat()</i>         | <i>wmemchr()</i>   |
| <i>wchr()</i>          | <i>wmemcmp()</i>   |
| <i>wscmp()</i>         | <i>wmemcpy()</i>   |
| <i>wscopy()</i>        | <i>wmemmove()</i>  |
| <i>wscspn()</i>        | <i>wmemset()</i>   |

See the “Caveats” section for the following functions for more information:

|                     |                               |
|---------------------|-------------------------------|
| <i>TraceEvent()</i> | <i>nanospin_ns()</i>          |
| <i>nanospin()</i>   | <i>nanospin_ns_to_count()</i> |

## Signal handlers

It’s safe to call the following functions from a signal handler:

|                           |                        |
|---------------------------|------------------------|
| <i>ChannelCreate()</i>    | <i>ClockAdjust_r()</i> |
| <i>ChannelCreate_r()</i>  | <i>ClockCycles()</i>   |
| <i>ChannelDestroy()</i>   | <i>ClockId()</i>       |
| <i>ChannelDestroy_r()</i> | <i>ClockId_r()</i>     |
| <i>ClockAdjust()</i>      | <i>ClockPeriod()</i>   |

|                                 |                             |
|---------------------------------|-----------------------------|
| <i>ClockPeriod_r()</i>          | <i>InterruptEnable()</i>    |
| <i>ClockTime()</i>              | <i>InterruptHookIdle()</i>  |
| <i>ClockTime_r()</i>            | <i>InterruptHookTrace()</i> |
| <i>ConnectAttach()</i>          | <i>InterruptLock()</i>      |
| <i>ConnectAttach_r()</i>        | <i>InterruptMask()</i>      |
| <i>ConnectClientInfo()</i>      | <i>InterruptUnlock()</i>    |
| <i>ConnectClientInfo_r()</i>    | <i>InterruptUnmask()</i>    |
| <i>ConnectDetach()</i>          | <i>InterruptWait()</i>      |
| <i>ConnectDetach_r()</i>        | <i>InterruptWait_r()</i>    |
| <i>ConnectFlags()</i>           | <i>MsgDeliverEvent()</i>    |
| <i>ConnectFlags_r()</i>         | <i>MsgDeliverEvent_r()</i>  |
| <i>ConnectServerInfo()</i>      | <i>MsgError()</i>           |
| <i>ConnectServerInfo_r()</i>    | <i>MsgError_r()</i>         |
| <i>DebugBreak()</i>             | <i>MsgInfo()</i>            |
| <i>DebugKDBreak()</i>           | <i>MsgInfo_r()</i>          |
| <i>DebugKDOOutput()</i>         | <i>MsgKeyData()</i>         |
| <i>ENDIAN_BE16()</i>            | <i>MsgKeyData_r()</i>       |
| <i>ENDIAN_BE32()</i>            | <i>MsgRead()</i>            |
| <i>ENDIAN_BE64()</i>            | <i>MsgRead_r()</i>          |
| <i>ENDIAN_LE16()</i>            | <i>MsgReady()</i>           |
| <i>ENDIAN_LE32()</i>            | <i>MsgReady_r()</i>         |
| <i>ENDIAN_LE64()</i>            | <i>MsgReceive()</i>         |
| <i>ENDIAN_RET16()</i>           | <i>MsgReceivePulse()</i>    |
| <i>ENDIAN_RET32()</i>           | <i>MsgReceivePulse_r()</i>  |
| <i>ENDIAN_RET64()</i>           | <i>MsgReceivePulsev()</i>   |
| <i>ENDIAN_SWAP16()</i>          | <i>MsgReceivePulsev_r()</i> |
| <i>ENDIAN_SWAP32()</i>          | <i>MsgReceive_r()</i>       |
| <i>ENDIAN_SWAP64()</i>          | <i>MsgReceivev()</i>        |
| <i>GETIOVBASE()</i>             | <i>MsgReceivev_r()</i>      |
| <i>GETIOVLEN()</i>              | <i>MsgReply()</i>           |
| <i>InterruptAttach()</i>        | <i>MsgReply_r()</i>         |
| <i>InterruptAttachEvent()</i>   | <i>MsgReplyv()</i>          |
| <i>InterruptAttachEvent_r()</i> | <i>MsgReplyv_r()</i>        |
| <i>InterruptAttach_r()</i>      | <i>MsgSend()</i>            |
| <i>InterruptDetach()</i>        | <i>MsgSendPulse()</i>       |
| <i>InterruptDetach_r()</i>      | <i>MsgSendPulse_r()</i>     |
| <i>InterruptDisable()</i>       | <i>MsgSend_r()</i>          |



|                            |                              |
|----------------------------|------------------------------|
| <i>MsgSendnc()</i>         | <i>SignalProcmask_r()</i>    |
| <i>MsgSendnc_r()</i>       | <i>SignalSuspend()</i>       |
| <i>MsgSendsv()</i>         | <i>SignalSuspend_r()</i>     |
| <i>MsgSendsv_r()</i>       | <i>SignalWaitinfo()</i>      |
| <i>MsgSendsvnc()</i>       | <i>SignalWaitinfo_r()</i>    |
| <i>MsgSendsvnc_r()</i>     | <i>SyncCondvarSignal()</i>   |
| <i>MsgSendv()</i>          | <i>SyncCondvarSignal_r()</i> |
| <i>MsgSendv_r()</i>        | <i>SyncCondvarWait()</i>     |
| <i>MsgSendvnc()</i>        | <i>SyncCondvarWait_r()</i>   |
| <i>MsgSendvnc_r()</i>      | <i>SyncCtl()</i>             |
| <i>MsgSendvs()</i>         | <i>SyncCtl_r()</i>           |
| <i>MsgSendvs_r()</i>       | <i>SyncDestroy()</i>         |
| <i>MsgSendvsnc()</i>       | <i>SyncDestroy_r()</i>       |
| <i>MsgSendvsnc_r()</i>     | <i>SyncMutexEvent()</i>      |
| <i>MsgVerifyEvent()</i>    | <i>SyncMutexEvent_r()</i>    |
| <i>MsgVerifyEvent_r()</i>  | <i>SyncMutexLock()</i>       |
| <i>MsgWrite()</i>          | <i>SyncMutexLock_r()</i>     |
| <i>MsgWrite_r()</i>        | <i>SyncMutexRevive()</i>     |
| <i>MsgWritev()</i>         | <i>SyncMutexRevive_r()</i>   |
| <i>MsgWritev_r()</i>       | <i>SyncMutexUnlock()</i>     |
| <i>ND_NODE_CMP()</i>       | <i>SyncMutexUnlock_r()</i>   |
| <i>SETIOV()</i>            | <i>SyncSemPost()</i>         |
| <i>SYSPAGE_CPU_ENTRY()</i> | <i>SyncSemPost_r()</i>       |
| <i>SYSPAGE_ENTRY()</i>     | <i>SyncSemWait()</i>         |
| <i>SchedGet()</i>          | <i>SyncSemWait_r()</i>       |
| <i>SchedGet_r()</i>        | <i>SyncTypeCreate()</i>      |
| <i>SchedInfo()</i>         | <i>SyncTypeCreate_r()</i>    |
| <i>SchedInfo_r()</i>       | <i>ThreadCancel()</i>        |
| <i>SchedSet()</i>          | <i>ThreadCancel_r()</i>      |
| <i>SchedSet_r()</i>        | <i>ThreadCreate()</i>        |
| <i>SchedYield()</i>        | <i>ThreadCreate_r()</i>      |
| <i>SchedYield_r()</i>      | <i>ThreadCtl()</i>           |
| <i>SignalAction()</i>      | <i>ThreadCtl_r()</i>         |
| <i>SignalAction_r()</i>    | <i>ThreadDestroy()</i>       |
| <i>SignalKill()</i>        | <i>ThreadDestroy_r()</i>     |
| <i>SignalKill_r()</i>      | <i>ThreadDetach()</i>        |
| <i>SignalProcmask()</i>    | <i>ThreadDetach_r()</i>      |

|                          |                              |
|--------------------------|------------------------------|
| <i>ThreadJoin()</i>      | <i>alloca()</i>              |
| <i>ThreadJoin_r()</i>    | <i>alphasort()</i>           |
| <i>TimerAlarm()</i>      | <i>asctime()</i>             |
| <i>TimerAlarm_r()</i>    | <i>asctime_r()</i>           |
| <i>TimerCreate()</i>     | <i>atoh()</i>                |
| <i>TimerCreate_r()</i>   | <i>atoi()</i>                |
| <i>TimerDestroy()</i>    | <i>atol()</i>                |
| <i>TimerDestroy_r()</i>  | <i>atoll()</i>               |
| <i>TimerInfo()</i>       | <i>atomic_add()</i>          |
| <i>TimerInfo_r()</i>     | <i>atomic_add_value()</i>    |
| <i>TimerSettime()</i>    | <i>atomic_clr()</i>          |
| <i>TimerSettime_r()</i>  | <i>atomic_clr_value()</i>    |
| <i>TimerTimeout()</i>    | <i>atomic_set()</i>          |
| <i>TimerTimeout_r()</i>  | <i>atomic_set_value()</i>    |
| <i>TraceEvent()</i>      | <i>atomic_sub()</i>          |
| <i>UNALIGNED_PUT16()</i> | <i>atomic_sub_value()</i>    |
| <i>UNALIGNED_PUT32()</i> | <i>atomic_toggle()</i>       |
| <i>UNALIGNED_PUT64()</i> | <i>atomic_toggle_value()</i> |
| <i>UNALIGNED_RET16()</i> | <i>basename()</i>            |
| <i>UNALIGNED_RET32()</i> | <i>bcmp()</i>                |
| <i>UNALIGNED_RET64()</i> | <i>bcopy()</i>               |
| <i>_RESMGR_NPARTS()</i>  | <i>bsearch()</i>             |
| <i>_RESMGR_PTR()</i>     | <i>btowc()</i>               |
| <i>_RESMGR_STATUS()</i>  | <i>bzero()</i>               |
| <i>_exit()</i>           | <i>cfgetispeed()</i>         |
| <i>_intr_v86()</i>       | <i>cfgetospeed()</i>         |
| <i>_sfree()</i>          | <i>cfgopen()</i>             |
| <i>abs()</i>             | <i>cfmakeraw()</i>           |
| <i>access()</i>          | <i>cfsetispeed()</i>         |
| <i>aio_cancel()</i>      | <i>cfsetospeed()</i>         |
| <i>aio_error()</i>       | <i>chdir()</i>               |
| <i>aio_fsync()</i>       | <i>chmod()</i>               |
| <i>aio_read()</i>        | <i>chown()</i>               |
| <i>aio_return()</i>      | <i>chsize()</i>              |
| <i>aio_suspend()</i>     | <i>clock()</i>               |
| <i>aio_write()</i>       | <i>clock_getcpuclockid()</i> |
| <i>alarm()</i>           | <i>clock_getres()</i>        |

|                           |                        |
|---------------------------|------------------------|
| <i>clock_gettime()</i>    | <i>fchown()</i>        |
| <i>clock_nanosleep()</i>  | <i>fcntl()</i>         |
| <i>clock_settime()</i>    | <i>fdatasync()</i>     |
| <i>close()</i>            | <i>ffs()</i>           |
| <i>confstr()</i>          | <i>fileno()</i>        |
| <i>creat()</i>            | <i>fink()</i>          |
| <i>creat64()</i>          | <i>flock()</i>         |
| <i>ctime()</i>            | <i>fnmatch()</i>       |
| <i>ctime_r()</i>          | <i>fork()</i>          |
| <i>daemon()</i>           | <i>forkpty()</i>       |
| <i>delay()</i>            | <i>fpathconf()</i>     |
| <i>devctl()</i>           | <i>fseek()</i>         |
| <i>dirname()</i>          | <i>fseeko()</i>        |
| <i>dispatch_block()</i>   | <i>fsetpos()</i>       |
| <i>dispatch_unblock()</i> | <i>fstat()</i>         |
| <i>div()</i>              | <i>fstat64()</i>       |
| <i>dn_comp()</i>          | <i>fstatvfs()</i>      |
| <i>dn_expand()</i>        | <i>fstatvfs64()</i>    |
| <i>ds_clear()</i>         | <i>fsync()</i>         |
| <i>ds_create()</i>        | <i>ftime()</i>         |
| <i>ds_deregister()</i>    | <i>ftruncate()</i>     |
| <i>ds_flags()</i>         | <i>ftruncate64()</i>   |
| <i>ds_get()</i>           | <i>ftrylockfile()</i>  |
| <i>ds_register()</i>      | <i>ftw()</i>           |
| <i>ds_set()</i>           | <i>ftw64()</i>         |
| <i>dup()</i>              | <i>futime()</i>        |
| <i>dup2()</i>             | <i>fwide()</i>         |
| <i>eaccess()</i>          | <i>gai_strerror()</i>  |
| <i>encrypt()</i>          | <i>getdomainname()</i> |
| <i>eof()</i>              | <i>getdtablesize()</i> |
| <i>err()</i>              | <i>getegid()</i>       |
| <i>errx()</i>             | <i>geteuid()</i>       |
| <i>execle()</i>           | <i>getgid()</i>        |
| <i>execve()</i>           | <i>getgrouplist()</i>  |
| <i>execvpe()</i>          | <i>getgroups()</i>     |
| <i>fcgopen()</i>          | <i>gethostname()</i>   |
| <i>fchmod()</i>           | <i>getitimer()</i>     |

*getpgid()*  
*getpgrp()*  
*getpid()*  
*getppid()*  
*getprio()*  
*getrlimit()*  
*getrlimit64()*  
*getrusage()*  
*getsubopt()*  
*gettimeofday()*  
*getuid()*  
*getw()*  
*getwd()*  
*glob()*  
*globfree()*  
*gmtime\_r()*  
*hsearch()*  
*hsterror()*  
*htonl()*  
*htons()*  
*hwi\_find\_item()*  
*hwi\_find\_tag()*  
*hwi\_off2tag()*  
*hwi\_tag2off()*  
*in16()*  
*in16s()*  
*in32()*  
*in32s()*  
*in8()*  
*in8s()*  
*inbe16()*  
*inbe32()*  
*index()*  
*inet6\_option\_\**  
*inet6\_rthdr\_\**  
*inet\_addr()*  
*inet\_aton()*  
*inet\_lnaof()*  
*inet\_makeaddr()*  
*inet\_netof()*  
*inet\_network()*  
*inet\_ntop()*  
*inet\_pton()*  
*inle16()*  
*inle32()*  
*iofdinfo()*  
*iofunc\_attr\_init()*  
*iofunc\_attr\_lock()*  
*iofunc\_attr\_trylock()*  
*iofunc\_attr\_unlock()*  
*iofunc\_check\_access()*  
*iofunc\_chmod()*  
*iofunc\_chmod\_default()*  
*iofunc\_chown()*  
*iofunc\_chown\_default()*  
*iofunc\_client\_info()*  
*iofunc\_close\_dup()*  
*iofunc\_close\_dup\_default()*  
*iofunc\_close\_ocb()*  
*iofunc\_close\_ocb\_default()*  
*iofunc\_devctl()*  
*iofunc\_devctl\_default()*  
*iofunc\_fdinfo()*  
*iofunc\_fdinfo\_default()*  
*iofunc\_func\_init()*  
*iofunc\_link()*  
*iofunc\_lock()*  
*iofunc\_lock\_alloc()*  
*iofunc\_lock\_default()*  
*iofunc\_lock\_free()*  
*iofunc\_lock\_ocb\_default()*  
*iofunc\_lseek()*  
*iofunc\_lseek\_default()*  
*iofunc\_mknod()*

|                                    |                      |
|------------------------------------|----------------------|
| <i>iofunc_mmap()</i>               | <i>isalpha()</i>     |
| <i>iofunc_mmap_default()</i>       | <i>isascii()</i>     |
| <i>iofunc_notify()</i>             | <i>iscntrl()</i>     |
| <i>iofunc_notify_remove()</i>      | <i>isdigit()</i>     |
| <i>iofunc_notify_trigger()</i>     | <i>isfdtype()</i>    |
| <i>iofunc_ocb_attach()</i>         | <i>isgraph()</i>     |
| <i>iofunc_ocb_calloc()</i>         | <i>islower()</i>     |
| <i>iofunc_ocb_detach()</i>         | <i>isprint()</i>     |
| <i>iofunc_ocb_free()</i>           | <i>ispunct()</i>     |
| <i>iofunc_open()</i>               | <i>isspace()</i>     |
| <i>iofunc_open_default()</i>       | <i>isupper()</i>     |
| <i>iofunc_openfd()</i>             | <i>iswalnum()</i>    |
| <i>iofunc_openfd_default()</i>     | <i>iswalpha()</i>    |
| <i>iofunc_pathconf()</i>           | <i>iswcntrl()</i>    |
| <i>iofunc_pathconf_default()</i>   | <i>iswctype()</i>    |
| <i>iofunc_read_default()</i>       | <i>iswdigit()</i>    |
| <i>iofunc_read_verify()</i>        | <i>iswgraph()</i>    |
| <i>iofunc_readlink()</i>           | <i>iswlower()</i>    |
| <i>iofunc_rename()</i>             | <i>iswprint()</i>    |
| <i>iofunc_space_verify()</i>       | <i>iswpunct()</i>    |
| <i>iofunc_stat()</i>               | <i>iswspace()</i>    |
| <i>iofunc_stat_default()</i>       | <i>iswupper()</i>    |
| <i>iofunc_sync()</i>               | <i>iswxdigit()</i>   |
| <i>iofunc_sync_default()</i>       | <i>isxdigit()</i>    |
| <i>iofunc_sync_verify()</i>        | <i>itoa()</i>        |
| <i>iofunc_time_update()</i>        | <i>kill()</i>        |
| <i>iofunc_unblock()</i>            | <i>killpg()</i>      |
| <i>iofunc_unblock_default()</i>    | <i>labs()</i>        |
| <i>iofunc_unlink()</i>             | <i>lchown()</i>      |
| <i>iofunc_unlock_ocb_default()</i> | <i>ldiv()</i>        |
| <i>iofunc_utime()</i>              | <i>lfind()</i>       |
| <i>iofunc_utime_default()</i>      | <i>link()</i>        |
| <i>iofunc_write_default()</i>      | <i>lio_listio()</i>  |
| <i>ionotify()</i>                  | <i>lltoa()</i>       |
| <i>ipsec_get_policylen()</i>       | <i>localtime_r()</i> |
| <i>ipsec_strerror()</i>            | <i>lockf()</i>       |
| <i>isalnum()</i>                   | <i>login_tty()</i>   |

|                             |                               |
|-----------------------------|-------------------------------|
| <i>longjmp()</i>            | <i>modem_open()</i>           |
| <i>lsearch()</i>            | <i>modem_write()</i>          |
| <i>lseek()</i>              | <i>mount()</i>                |
| <i>lseek64()</i>            | <i>mprotect()</i>             |
| <i>lstat()</i>              | <i>mq_timedreceive()</i>      |
| <i>lstat64()</i>            | <i>mq_timedsend()</i>         |
| <i>ltoa()</i>               | <i>msync()</i>                |
| <i>max()</i>                | <i>munmap()</i>               |
| <i>mblen()</i>              | <i>munmap_device_io()</i>     |
| <i>mbrlen()</i>             | <i>munmap_device_memory()</i> |
| <i>mbrtowc()</i>            | <i>name_close()</i>           |
| <i>mbsinit()</i>            | <i>name_open()</i>            |
| <i>mbsrtowcs()</i>          | <i>nanospin()</i>             |
| <i>mbstowcs()</i>           | <i>nanospin_calibrate()</i>   |
| <i>mbtowc()</i>             | <i>nanospin_count()</i>       |
| <i>mem_offset()</i>         | <i>nanospin_ns()</i>          |
| <i>mem_offset64()</i>       | <i>nanospin_ns_to_count()</i> |
| <i>memalign()</i>           | <i>nap()</i>                  |
| <i>memccpy()</i>            | <i>napms()</i>                |
| <i>memchr()</i>             | <i>nbaconnect_result()</i>    |
| <i>memcmp()</i>             | <i>netmgr_ndtostr()</i>       |
| <i>memcpy()</i>             | <i>netmgr_remote_nd()</i>     |
| <i>memcpyv()</i>            | <i>netmgr_strtond()</i>       |
| <i>memicmp()</i>            | <i>nftw()</i>                 |
| <i>memmove()</i>            | <i>nftw64()</i>               |
| <i>memset()</i>             | <i>nice()</i>                 |
| <i>min()</i>                | <i>nsec2timespec()</i>        |
| <i>mkdir()</i>              | <i>ntohl()</i>                |
| <i>mkfifo()</i>             | <i>ntohs()</i>                |
| <i>mknod()</i>              | <i>offsetof()</i>             |
| <i>mkstemp()</i>            | <i>open()</i>                 |
| <i>mktemp()</i>             | <i>open64()</i>               |
| <i>mktime()</i>             | <i>openfd()</i>               |
| <i>mmap()</i>               | <i>openpty()</i>              |
| <i>mmap64()</i>             | <i>out16()</i>                |
| <i>mmap_device_io()</i>     | <i>out16s()</i>               |
| <i>mmap_device_memory()</i> | <i>out32()</i>                |

|                                  |                                       |
|----------------------------------|---------------------------------------|
| <i>out32s()</i>                  | <i>pci_write_config8()</i>            |
| <i>out8()</i>                    | <i>pipe()</i>                         |
| <i>out8s()</i>                   | <i>posix_mem_offset()</i>             |
| <i>outbe16()</i>                 | <i>posix_mem_offset64()</i>           |
| <i>outbe32()</i>                 | <i>posix_memalign()</i>               |
| <i>outle16()</i>                 | <i>pread()</i>                        |
| <i>outle32()</i>                 | <i>pread64()</i>                      |
| <i>pathconf()</i>                | <i>procmgr_daemon()</i>               |
| <i>pathfind()</i>                | <i>procmgr_event_notify()</i>         |
| <i>pathfind_r()</i>              | <i>procmgr_event_trigger()</i>        |
| <i>pathmgr_symlink()</i>         | <i>procmgr_guardian()</i>             |
| <i>pathmgr_unlink()</i>          | <i>pthread_abort()</i>                |
| <i>pause()</i>                   | <i>pthread_atfork()</i>               |
| <i>pccard_arm()</i>              | <i>pthread_attr_destroy()</i>         |
| <i>pccard_attach()</i>           | <i>pthread_attr_getdetachstate()</i>  |
| <i>pccard_detach()</i>           | <i>pthread_attr_getguardsize()</i>    |
| <i>pccard_info()</i>             | <i>pthread_attr_getinheritsched()</i> |
| <i>pccard_lock()</i>             | <i>pthread_attr_getschedparam()</i>   |
| <i>pccard_raw_read()</i>         | <i>pthread_attr_getschedpolicy()</i>  |
| <i>pccard_unlock()</i>           | <i>pthread_attr_getscope()</i>        |
| <i>pci_attach()</i>              | <i>pthread_attr_getstackaddr()</i>    |
| <i>pci_attach_device()</i>       | <i>pthread_attr_getstacklazy()</i>    |
| <i>pci_detach()</i>              | <i>pthread_attr_getstacksize()</i>    |
| <i>pci_detach_device()</i>       | <i>pthread_attr_init()</i>            |
| <i>pci_find_class()</i>          | <i>pthread_attr_setdetachstate()</i>  |
| <i>pci_find_device()</i>         | <i>pthread_attr_setguardsize()</i>    |
| <i>pci_irq_routing_options()</i> | <i>pthread_attr_setinheritsched()</i> |
| <i>pci_map_irq()</i>             | <i>pthread_attr_setschedparam()</i>   |
| <i>pci_present()</i>             | <i>pthread_attr_setschedpolicy()</i>  |
| <i>pci_read_config()</i>         | <i>pthread_attr_setscope()</i>        |
| <i>pci_read_config16()</i>       | <i>pthread_attr_setstackaddr()</i>    |
| <i>pci_read_config32()</i>       | <i>pthread_attr_setstacklazy()</i>    |
| <i>pci_read_config8()</i>        | <i>pthread_attr_setstacksize()</i>    |
| <i>pci_rescan_bus()</i>          | <i>pthread_barrier_destroy()</i>      |
| <i>pci_write_config()</i>        | <i>pthread_barrier_init()</i>         |
| <i>pci_write_config16()</i>      | <i>pthread_barrier_wait()</i>         |
| <i>pci_write_config32()</i>      | <i>pthread_barrierattr_destroy()</i>  |

*pthread\_barrierattr\_getpshared()*  
*pthread\_barrierattr\_init()*  
*pthread\_barrierattr\_setpshared()*  
*pthread\_cancel()*  
*pthread\_cleanup\_pop()*  
*pthread\_cleanup\_push()*  
*pthread\_cond\_broadcast()*  
*pthread\_cond\_destroy()*  
*pthread\_cond\_init()*  
*pthread\_cond\_signal()*  
*pthread\_cond\_timedwait()*  
*pthread\_cond\_wait()*  
*pthread\_condattr\_destroy()*  
*pthread\_condattr\_getclock()*  
*pthread\_condattr\_getpshared()*  
*pthread\_condattr\_init()*  
*pthread\_condattr\_setclock()*  
*pthread\_condattr\_setpshared()*  
*pthread\_create()*  
*pthread\_detach()*  
*pthread\_equal()*  
*pthread\_exit()*  
*pthread\_getconcurrency()*  
*pthread\_getcpuclockid()*  
*pthread\_getschedparam()*  
*pthread\_getspecific()*  
*pthread\_join()*  
*pthread\_key\_delete()*  
*pthread\_kill()*  
*pthread\_mutex\_destroy()*  
*pthread\_mutex\_getprioceiling()*  
*pthread\_mutex\_init()*  
*pthread\_mutex\_lock()*  
*pthread\_mutex\_setprioceiling()*  
*pthread\_mutex\_timedlock()*  
*pthread\_mutex\_trylock()*  
*pthread\_mutex\_unlock()*  
*pthread\_mutexattr\_destroy()*  
*pthread\_mutexattr\_getprioceiling()*  
*pthread\_mutexattr\_getprotocol()*  
*pthread\_mutexattr\_getpshared()*  
*pthread\_mutexattr\_getrecursive()*  
*pthread\_mutexattr\_gettype()*  
*pthread\_mutexattr\_init()*  
*pthread\_mutexattr\_setprioceiling()*  
*pthread\_mutexattr\_setprotocol()*  
*pthread\_mutexattr\_setpshared()*  
*pthread\_mutexattr\_setrecursive()*  
*pthread\_mutexattr\_settype()*  
*pthread\_once()*  
*pthread\_rwlock\_destroy()*  
*pthread\_rwlock\_init()*  
*pthread\_rwlock\_rdlock()*  
*pthread\_rwlock\_timedrdlock()*  
*pthread\_rwlock\_timedwrlock()*  
*pthread\_rwlock\_tryrdlock()*  
*pthread\_rwlock\_trywrlock()*  
*pthread\_rwlock\_unlock()*  
*pthread\_rwlock\_wrlock()*  
*pthread\_rwlockattr\_destroy()*  
*pthread\_rwlockattr\_getpshared()*  
*pthread\_rwlockattr\_init()*  
*pthread\_rwlockattr\_setpshared()*  
*pthread\_self()*  
*pthread\_setcancelstate()*  
*pthread\_setcanceltype()*  
*pthread\_setconcurrency()*  
*pthread\_setschedparam()*  
*pthread\_sigmask()*  
*pthread\_sleepon\_broadcast()*  
*pthread\_sleepon\_lock()*  
*pthread\_sleepon\_signal()*  
*pthread\_sleepon\_timedwait()*  
*pthread\_sleepon\_unlock()*



|                               |                                    |
|-------------------------------|------------------------------------|
| <i>pthread_sleepon_wait()</i> | <i>rsrddbmgr_create()</i>          |
| <i>pthread_spin_destroy()</i> | <i>rsrddbmgr_destroy()</i>         |
| <i>pthread_spin_init()</i>    | <i>rsrddbmgr_detach()</i>          |
| <i>pthread_spin_lock()</i>    | <i>rsrddbmgr_devno_attach()</i>    |
| <i>pthread_spin_trylock()</i> | <i>rsrddbmgr_devno_detach()</i>    |
| <i>pthread_spin_unlock()</i>  | <i>rsrddbmgr_query()</i>           |
| <i>pthread_testcancel()</i>   | <i>scandir()</i>                   |
| <i>pthread_timedjoin()</i>    | <i>sched_get_priority_adjust()</i> |
| <i>putw()</i>                 | <i>sched_get_priority_max()</i>    |
| <i>pwrite()</i>               | <i>sched_get_priority_min()</i>    |
| <i>pwrite64()</i>             | <i>sched_getparam()</i>            |
| <i>qnx_crypt()</i>            | <i>sched_getscheduler()</i>        |
| <i>raise()</i>                | <i>sched_rr_get_interval()</i>     |
| <i>rand()</i>                 | <i>sched_setparam()</i>            |
| <i>rand_r()</i>               | <i>sched_setscheduler()</i>        |
| <i>random()</i>               | <i>sched_yield()</i>               |
| <i>rdchk()</i>                | <i>sem_close()</i>                 |
| <i>re_comp()</i>              | <i>sem_destroy()</i>               |
| <i>re_exec()</i>              | <i>sem_getvalue()</i>              |
| <i>read()</i>                 | <i>sem_open()</i>                  |
| <i>readblock()</i>            | <i>sem_post()</i>                  |
| <i>readcond()</i>             | <i>sem_timedwait()</i>             |
| <i>readdir_r()</i>            | <i>sem_trywait()</i>               |
| <i>readlink()</i>             | <i>sem_unlink()</i>                |
| <i>readv()</i>                | <i>sem_wait()</i>                  |
| <i>realpath()</i>             | <i>setdomainname()</i>             |
| <i>regerror()</i>             | <i>setegid()</i>                   |
| <i>rename()</i>               | <i>seteuid()</i>                   |
| <i>resmgr_msgread()</i>       | <i>setgid()</i>                    |
| <i>resmgr_msgreadv()</i>      | <i>sethostname()</i>               |
| <i>resmgr_msgwrite()</i>      | <i>setitimer()</i>                 |
| <i>resmgr_msgwritev()</i>     | <i>setjmp()</i>                    |
| <i>resmgr_pathname()</i>      | <i>setpgid()</i>                   |
| <i>rewinddir()</i>            | <i>setpgrp()</i>                   |
| <i>rindex()</i>               | <i>setprio()</i>                   |
| <i>rmdir()</i>                | <i>setregid()</i>                  |
| <i>rsrddbmgr_attach()</i>     | <i>setreuid()</i>                  |

|                       |                      |
|-----------------------|----------------------|
| <i>setrlimit()</i>    | <i>sprintf()</i>     |
| <i>setrlimit64()</i>  | <i>srand()</i>       |
| <i>setsid()</i>       | <i>srand48()</i>     |
| <i>settimeofday()</i> | <i>sscanf()</i>      |
| <i>setuid()</i>       | <i>stat()</i>        |
| <i>setutent()</i>     | <i>stat64()</i>      |
| <i>shm_ctl()</i>      | <i>statvfs()</i>     |
| <i>shm_open()</i>     | <i>statvfs64()</i>   |
| <i>shm_unlink()</i>   | <i>straddstr()</i>   |
| <i>sigaction()</i>    | <i>strcasecmp()</i>  |
| <i>sigaddset()</i>    | <i>strcat()</i>      |
| <i>sigblock()</i>     | <i>strchr()</i>      |
| <i>sigdelset()</i>    | <i>strcmp()</i>      |
| <i>sigemptyset()</i>  | <i>strcmpi()</i>     |
| <i>sigfillset()</i>   | <i>strcoll()</i>     |
| <i>sigismember()</i>  | <i>strcpy()</i>      |
| <i>siglongjmp()</i>   | <i>strcspn()</i>     |
| <i>sigmask()</i>      | <i>strerror()</i>    |
| <i>signal()</i>       | <i>strftime()</i>    |
| <i>sigpause()</i>     | <i>stricmp()</i>     |
| <i>sigpending()</i>   | <i>strlen()</i>      |
| <i>sigprocmask()</i>  | <i>strlwr()</i>      |
| <i>sigqueue()</i>     | <i>strncasecmp()</i> |
| <i>sigsetjmp()</i>    | <i>strncat()</i>     |
| <i>sigsetmask()</i>   | <i>strncmp()</i>     |
| <i>sigsuspend()</i>   | <i>strncpy()</i>     |
| <i>sigtimedwait()</i> | <i>strnicmp()</i>    |
| <i>sigunblock()</i>   | <i>strnset()</i>     |
| <i>sigwait()</i>      | <i>strpbrk()</i>     |
| <i>sigwaitinfo()</i>  | <i>strrchr()</i>     |
| <i>sleep()</i>        | <i>strrev()</i>      |
| <i>slogb()</i>        | <i>strsep()</i>      |
| <i>slogf()</i>        | <i>strset()</i>      |
| <i>slogi()</i>        | <i>strsignal()</i>   |
| <i>snprintf()</i>     | <i>strspn()</i>      |
| <i>sopenfd()</i>      | <i>strstr()</i>      |
| <i>spawn()</i>        | <i>strtod()</i>      |

*strtoimax()*  
*strtok\_r()*  
*strtol()*  
*strtoll()*  
*strtoul()*  
*strtoull()*  
*strtoumax()*  
*strupr()*  
*strxfrm()*  
*swab()*  
*swprintf()*  
*swscanf()*  
*symlink()*  
*sync()*  
*sysconf()*  
*sysmgr\_reboot()*  
*tcdrain()*  
*tcdropline()*  
*tcflow()*  
*tcflush()*  
*tcgetattr()*  
*tcgetpgrp()*  
*tcgetsid()*  
*tcgetsize()*  
*tcinject()*  
*tcischars()*  
*tcsendbreak()*  
*tcsetattr()*  
*tcsetpgrp()*  
*tcsetsize()*  
*tell()*  
*tell64()*  
*time()*  
*timer\_create()*  
*timer\_delete()*  
*timer\_getexpstatus()*  
*timer\_getoverrun()*  
*timer\_gettime()*  
*timer\_settime()*  
*timer\_timeout()*  
*timer\_timeout\_r()*  
*times()*  
*timespec2nsec()*  
*tolower()*  
*toupper()*  
*towctrans()*  
*towlower()*  
*towupper()*  
*truncate()*  
*ttyname\_r()*  
*ualarm()*  
*ulltoa()*  
*ultoa()*  
*umask()*  
*uname()*  
*unlink()*  
*unsetenv()*  
*usleep()*  
*utime()*  
*utimes()*  
*utmpname()*  
*utoa()*  
*va\_arg()*  
*va\_copy()*  
*va\_end()*  
*va\_start()*  
*valloc()*  
*verr()*  
*verrx()*  
*vslogf()*  
*vwarn()*  
*vwarnx()*  
*wait()*  
*wait3()*

|                    |                     |
|--------------------|---------------------|
| <i>wait4()</i>     | <i>wcstof()</i>     |
| <i>waitid()</i>    | <i>wcstoimax()</i>  |
| <i>waitpid()</i>   | <i>wcstok()</i>     |
| <i>warn()</i>      | <i>wcstol()</i>     |
| <i>warnx()</i>     | <i>wcstold()</i>    |
| <i>wrtomb()</i>    | <i>wcstoll()</i>    |
| <i>wscat()</i>     | <i>wcstombs()</i>   |
| <i>wchr()</i>      | <i>wcstoul()</i>    |
| <i>wscmp()</i>     | <i>wcstoull()</i>   |
| <i>wscoll()</i>    | <i>wcstoumax()</i>  |
| <i>wscopy()</i>    | <i>wctob()</i>      |
| <i>wscspn()</i>    | <i>wctomb()</i>     |
| <i>wscxfrm()</i>   | <i>wctrans()</i>    |
| <i>wcsftime()</i>  | <i>wctype()</i>     |
| <i>wcslen()</i>    | <i>wmemchr()</i>    |
| <i>wcscat()</i>    | <i>wmemcmp()</i>    |
| <i>wcscmp()</i>    | <i>wmemcpy()</i>    |
| <i>wcscpy()</i>    | <i>wmemmove()</i>   |
| <i>wcspbrk()</i>   | <i>wmemset()</i>    |
| <i>wcsrchr()</i>   | <i>wordexp()</i>    |
| <i>wcsrtombs()</i> | <i>wordfree()</i>   |
| <i>wcsspn()</i>    | <i>write()</i>      |
| <i>wcsstr()</i>    | <i>writeblock()</i> |
| <i>wcstod()</i>    | <i>writev()</i>     |

See the “Caveats” section for the following functions for more information:

|                       |                    |
|-----------------------|--------------------|
| <i>abort()</i>        | <i>vsprintf()</i>  |
| <i>modem_read()</i>   | <i>vsscanf()</i>   |
| <i>modem_script()</i> | <i>vswprintf()</i> |
| <i>vsnprintf()</i>    | <i>vswscanf()</i>  |

## Multithreaded programs



**CAUTION:** It *isn't* safe to call these functions from a multithreaded program.

|                           |                           |
|---------------------------|---------------------------|
| <i>Raccept()</i>          | <i>getnameinfo()</i>      |
| <i>Rbind()</i>            | <i>getnetbyaddr()</i>     |
| <i>Rconnect()</i>         | <i>getnetbyname()</i>     |
| <i>Rgetsockname()</i>     | <i>getnetent()</i>        |
| <i>Rlisten()</i>          | <i>getopt()</i>           |
| <i>Rrcmd()</i>            | <i>getpass()</i>          |
| <i>Rselect()</i>          | <i>getprotobyname()</i>   |
| <i>SOCKSinit()</i>        | <i>getprotobynumber()</i> |
| <i>bindresvport()</i>     | <i>getprotoent()</i>      |
| <i>crypt()</i>            | <i>getpwent()</i>         |
| <i>daemon()</i>           | <i>getpwnam()</i>         |
| <i>drand48()</i>          | <i>getpwuid()</i>         |
| <i>endgrent()</i>         | <i>getservbyname()</i>    |
| <i>endhostent()</i>       | <i>getservbyport()</i>    |
| <i>endnetent()</i>        | <i>getservent()</i>       |
| <i>endprotoent()</i>      | <i>gmtime()</i>           |
| <i>endpwent()</i>         | <i>herror()</i>           |
| <i>endservent()</i>       | <i>inet_ntoa()</i>        |
| <i>endspent()</i>         | <i>initgroups()</i>       |
| <i>endutent()</i>         | <i>initstate()</i>        |
| <i>fgetspent()</i>        | <i>input_line()</i>       |
| <i>getaddrinfo()</i>      | <i>ioctl()</i>            |
| <i>getc_unlocked()</i>    | <i>lcong48()</i>          |
| <i>getchar_unlocked()</i> | <i>localtime()</i>        |
| <i>getenv()</i>           | <i>lrand48()</i>          |
| <i>getgrent()</i>         | <i>lstat()</i>            |
| <i>getgrgid()</i>         | <i>lstat64()</i>          |
| <i>getgrnam()</i>         | <i>mrand48()</i>          |
| <i>gethostbyaddr()</i>    | <i>pclose()</i>           |
| <i>gethostbyname()</i>    | <i>popen()</i>            |
| <i>gethostbyname2()</i>   | <i>putc_unlocked()</i>    |
| <i>gethostent()</i>       | <i>putchar_unlocked()</i> |
| <i>getlogin()</i>         | <i>putenv()</i>           |

|                                |                           |
|--------------------------------|---------------------------|
| <i>putspent()</i>              | <i>setpwent()</i>         |
| <i>rand()</i>                  | <i>setservent()</i>       |
| <i>random()</i>                | <i>setspent()</i>         |
| <i>rcmd()</i>                  | <i>setstate()</i>         |
| <i>read_main_config_file()</i> | <i>sigblock()</i>         |
| <i>readdir()</i>               | <i>snmp_close()</i>       |
| <i>res_init()</i>              | <i>snmp_free_pdu()</i>    |
| <i>res_mkquery()</i>           | <i>snmp_open()</i>        |
| <i>res_query()</i>             | <i>snmp_pdu_create()</i>  |
| <i>res_querydomain()</i>       | <i>snmp_read()</i>        |
| <i>res_search()</i>            | <i>snmp_select_info()</i> |
| <i>res_send()</i>              | <i>snmp_send()</i>        |
| <i>ruserok()</i>               | <i>snmp_timeout()</i>     |
| <i>seekdir()</i>               | <i>socketmark()</i>       |
| <i>select()</i>                | <i>srandom()</i>          |
| <i>setenv()</i>                | <i>strtok()</i>           |
| <i>setgrent()</i>              | <i>syslog()</i>           |
| <i>setgroups()</i>             | <i>telldir()</i>          |
| <i>sethostent()</i>            | <i>tempnam()</i>          |
| <i>setkey()</i>                | <i>ttyname()</i>          |
| <i>setlogmask()</i>            | <i>vfork()</i>            |
| <i>setnetent()</i>             | <i>vsyslog()</i>          |
| <i>setprotoent()</i>           | <i>wscoll()</i>           |

See the “Caveats” section for the following functions for more information:

|                     |                       |
|---------------------|-----------------------|
| <i>ctermid()</i>    | <i>modem_script()</i> |
| <i>modem_read()</i> | <i>tmpnam()</i>       |

## ***Glossary***

---





## A20 gate

On x86-based systems, a hardware component that forces the A20 address line on the bus to zero, regardless of the actual setting of the A20 address line on the processor. This component is in place to support legacy systems, but the QNX Neutrino OS doesn't require any such hardware. Note that some processors, such as the 386EX, have the A20 gate hardware built right into the processor itself — our IPL will disable the A20 gate as soon as possible after startup.

## adaptive

Scheduling algorithm whereby a thread's priority is decayed by 1. See also **FIFO**, **round robin**, and **sporadic**.

## atomic

Of or relating to atoms. :-)

In operating systems, this refers to the requirement that an operation, or sequence of operations, be considered *indivisible*. For example, a thread may need to move a file position to a given location and read data. These operations must be performed in an atomic manner; otherwise, another thread could preempt the original thread and move the file position to a different location, thus causing the original thread to read data from the second thread's position.

## attributes structure

Structure containing information used on a per-resource basis (as opposed to the **OCB**, which is used on a per-open basis).

This structure is also known as a **handle**. The structure definition is fixed (`iofunc_attr_t`), but may be extended. See also **mount structure**.

## bank-switched

A term indicating that a certain memory component (usually the device holding an **image**) isn't entirely addressable by the processor. In this case, a hardware component manifests a small portion (or "window") of the device onto the processor's address bus. Special

commands have to be issued to the hardware to move the window to different locations in the device. See also **linearly mapped**.

### base layer calls

Convenient set of library calls for writing resource managers. These calls all start with *resmgr\_\**(*resmgr\_pathname\_attach()*). Note that while some base layer calls are unavoidable (e.g. *resmgr\_pathname\_attach()*), we recommend that you use the **POSIX layer calls** where possible.

### BIOS/ROM Monitor extension signature

A certain sequence of bytes indicating to the BIOS or ROM Monitor that the device is to be considered an “extension” to the BIOS or ROM Monitor — control is to be transferred to the device by the BIOS or ROM Monitor, with the expectation that the device will perform additional initializations.

On the x86 architecture, the two bytes 0x55 and 0xAA must be present (in that order) as the first two bytes in the device, with control being transferred to offset 0x0003.

### block-integral

The requirement that data be transferred such that individual structure components are transferred in their entirety — no partial structure component transfers are allowed.

In a resource manager, directory data must be returned to a client as **block-integral** data. This means that only complete **struct dirent** structures can be returned — it’s inappropriate to return partial structures, assuming that the next `IO_READ` request will “pick up” where the previous one left off.

### bootable

An image can be either bootable or **nonbootable**. A bootable image is one that contains the startup code that the IPL can transfer control to.

## **bootfile**

The part of an OS image that runs the **startup code** and the Neutrino microkernel.

## **budget**

In **sporadic** scheduling, the amount of time a thread is permitted to execute at its normal priority before being dropped to its low priority.

## **buildfile**

A text file containing instructions for **mki fs** specifying the contents and other details of an **image**, or for **mke fs** specifying the contents and other details of an embedded filesystem image.

## **canonical mode**

Also called edited mode or “cooked” mode. In this mode the character device library performs line-editing operations on each received character. Only when a line is “completely entered” — typically when a carriage return (CR) is received — will the line of data be made available to application processes. Contrast **raw mode**.

## **channel**

A kernel object used with message passing.

In QNX Neutrino, message passing is directed towards a **connection** (made to a channel); threads can receive messages from channels. A thread that wishes to receive messages creates a channel (using *ChannelCreate()*), and then receives messages from that channel (using *MsgReceive()*). Another thread that wishes to send a message to the first thread must make a connection to that channel by “attaching” to the channel (using *ConnectAttach()*) and then sending data (using *MsgSend()*).

## **CIFS**

Common Internet File System (aka SMB) — a protocol that allows a client workstation to perform transparent file access over a network to a Windows 95/98/NT server. Client file access calls are converted to

CIFS protocol requests and are sent to the server over the network. The server receives the request, performs the actual filesystem operation, and sends a response back to the client.

## CIS

Card Information Structure — a data block that maintains information about flash configuration. The CIS description includes the types of memory devices in the regions, the physical geometry of these devices, and the partitions located on the flash.

## combine message

A resource manager message that consists of two or more messages. The messages are constructed as combine messages by the client's C library (e.g. *stat()*, *readblock()*), and then handled as individual messages by the resource manager.

The purpose of combine messages is to conserve network bandwidth and/or to provide support for atomic operations. See also **connect message** and **I/O message**.

## connect message

In a resource manager, a message issued by the client to perform an operation based on a pathname (e.g. an **io\_open** message). Depending on the type of connect message sent, a context block (see **OCB**) may be associated with the request and will be passed to subsequent I/O messages. See also **combine message** and **I/O message**.

## connection

A kernel object used with message passing.

Connections are created by client threads to “connect” to the channels made available by servers. Once connections are established, clients can *MsgSendv()* messages over them. If a number of threads in a process all attach to the same channel, then the one connection is shared among all the threads. Channels and connections are identified within a process by a small integer.

The key thing to note is that connections and file descriptors (**FD**) are one and the same object. See also **channel** and **FD**.

### **context**

Information retained between invocations of functionality.

When using a resource manager, the client sets up an association or **context** within the resource manager by issuing an *open()* call and getting back a file descriptor. The resource manager is responsible for storing the information required by the context (see **OCB**). When the client issues further file-descriptor based messages, the resource manager uses the OCB to determine the context for interpretation of the client's messages.

### **cooked mode**

See **canonical mode**.

### **core dump**

A file describing the state of a process that terminated abnormally.

### **critical section**

A code passage that *must* be executed “serially” (i.e. by only one thread at a time). The simplest form of critical section enforcement is via a **mutex**.

### **deadlock**

A condition in which one or more threads are unable to continue due to resource contention. A common form of deadlock can occur when one thread sends a message to another, while the other thread sends a message to the first. Both threads are now waiting for each other to reply to the message. Deadlock can be avoided by good design practices or massive kludges — we recommend the good design approach.

## device driver

A process that allows the OS and application programs to make use of the underlying hardware in a generic way (e.g. a disk drive, a network interface). Unlike OSs that require device drivers to be tightly bound into the OS itself, device drivers for QNX Neutrino are standard processes that can be started and stopped dynamically. As a result, adding device drivers doesn't affect any other part of the OS — drivers can be developed and debugged like any other application. Also, device drivers are in their own protected address space, so a bug in a device driver won't cause the entire OS to shut down.

## DNS

Domain Name Service — an Internet protocol used to convert ASCII domain names into IP addresses. In QNX native networking, **dns** is one of **Qnet**'s builtin resolvers.

## dynamic bootfile

An OS image built on the fly. Contrast **static bootfile**.

## dynamic linking

The process whereby you link your modules in such a way that the Process Manager will link them to the library modules before your program runs. The word “dynamic” here means that the association between your program and the library modules that it uses is done *at load time*, not at linktime. Contrast **static linking**. See also **runtime loading**.

## edge-sensitive

One of two ways in which a **PIC** (Programmable Interrupt Controller) can be programmed to respond to interrupts. In edge-sensitive mode, the interrupt is “noticed” upon a transition to/from the rising/falling edge of a pulse. Contrast **level-sensitive**.

## edited mode

See **canonical mode**.

## EOI

End Of Interrupt — a command that the OS sends to the PIC after processing all Interrupt Service Routines (ISR) for that particular interrupt source so that the PIC can reset the processor's In Service Register. See also **PIC** and **ISR**.

## EPROM

Erasable Programmable Read-Only Memory — a memory technology that allows the device to be programmed (typically with higher-than-operating voltages, e.g. 12V), with the characteristic that any bit (or bits) may be individually programmed from a 1 state to a 0 state. To change a bit from a 0 state into a 1 state can only be accomplished by erasing the *entire* device, setting *all* of the bits to a 1 state. Erasing is accomplished by shining an ultraviolet light through the erase window of the device for a fixed period of time (typically 10-20 minutes). The device is further characterized by having a limited number of erase cycles (typically 10e5 - 10e6). Contrast **flash** and **RAM**.

## event

A notification scheme used to inform a thread that a particular condition has occurred. Events can be signals or pulses in the general case; they can also be unblocking events or interrupt events in the case of kernel timeouts and interrupt service routines. An event is delivered by a thread, a timer, the kernel, or an interrupt service routine when appropriate to the requestor of the event.

## FD

File Descriptor — a client must open a file descriptor to a resource manager via the *open()* function call. The file descriptor then serves as a handle for the client to use in subsequent messages. Note that a file descriptor is the exact same object as a connection ID (*coid*, returned by *ConnectAttach()*).

## **FIFO**

First In First Out — a scheduling algorithm whereby a thread is able to consume CPU at its priority level without bounds. See also **adaptive**, **round robin**, and **sporadic**.

## **flash memory**

A memory technology similar in characteristics to **EPROM** memory, with the exception that erasing is performed electrically instead of via ultraviolet light, and, depending upon the organization of the flash memory device, erasing may be accomplished in blocks (typically 64k bytes at a time) instead of the entire device. Contrast **EPROM** and **RAM**.

## **FQNN**

Fully Qualified NodeName — a unique name that identifies a QNX Neutrino node on a network. The FQNN consists of the nodename plus the node domain tacked together.

## **garbage collection**

Aka space reclamation, the process whereby a filesystem manager recovers the space occupied by deleted files and directories.

## **HA**

High Availability — in telecommunications and other industries, HA describes a system's ability to remain up and running without interruption for extended periods of time.

## **handle**

A pointer that the resource manager base library binds to the pathname registered via *resmgr\_attach()*. This handle is typically used to associate some kind of per-device information. Note that if you use the *iofunc\_\**(*)* **POSIX layer calls**, you must use a particular *type* of handle — in this case called an **attributes structure**.



## image

In the context of embedded QNX Neutrino systems, an “image” can mean either a structure that contains files (i.e. an OS image) or a structure that can be used in a read-only, read/write, or read/write/reclaim FFS-2-compatible filesystem (i.e. a flash filesystem image).

## interrupt

An event (usually caused by hardware) that interrupts whatever the processor was doing and asks it do something else. The hardware will generate an interrupt whenever it has reached some state where software intervention is required.

## interrupt handler

See **ISR**.

## interrupt latency

The amount of elapsed time between the generation of a hardware interrupt and the first instruction executed by the relevant interrupt service routine. Also designated as “ $T_{ij}$ ”. Contrast **scheduling latency**.

## interrupt service routine

See **ISR**.

## interrupt service thread

A thread that is responsible for performing thread-level servicing of an interrupt.

Since an **ISR** can call only a very limited number of functions, and since the amount of time spent in an ISR should be kept to a minimum, generally the bulk of the interrupt servicing work should be done by a thread. The thread attaches the interrupt (via *InterruptAttach()* or *InterruptAttachEvent()*) and then blocks (via *InterruptWait()*), waiting for the ISR to tell it to do something (by returning an event of type SIGEV\_INTR). To aid in minimizing

**scheduling latency**, the interrupt service thread should raise its priority appropriately.

### **I/O message**

A message that relies on an existing binding between the client and the resource manager. For example, an `_IO_READ` message depends on the client's having previously established an association (or **context**) with the resource manager by issuing an `open()` and getting back a file descriptor. See also **connect message**, **context**, **combine message**, and **message**.

### **I/O privity**

A particular privilege, that, if enabled for a given thread, allows the thread to perform I/O instructions (such as the x86 assembler `in` and `out` instructions). By default, I/O privity is disabled, because a program with it enabled can wreak havoc on a system. To enable I/O privity, the thread must be running as `root`, and call `ThreadCtl()`.

### **IPC**

Interprocess Communication — the ability for two processes (or threads) to communicate. QNX Neutrino offers several forms of IPC, most notably native messaging (synchronous, client/server relationship), POSIX message queues and pipes (asynchronous), as well as signals.

### **IPL**

Initial Program Loader — the software component that either takes control at the processor's reset vector (e.g. location `0xFFFFFFFF0` on the x86), or is a BIOS extension. This component is responsible for setting up the machine into a usable state, such that the startup program can then perform further initializations. The IPL is written in assembler and C. See also **BIOS extension signature** and **startup code**.

## **IRQ**

Interrupt Request — a hardware request line asserted by a peripheral to indicate that it requires servicing by software. The IRQ is handled by the **PIC**, which then interrupts the processor, usually causing the processor to execute an **Interrupt Service Routine (ISR)**.

## **ISR**

Interrupt Service Routine — a routine responsible for servicing hardware (e.g. reading and/or writing some device ports), for updating some data structures shared between the ISR and the thread(s) running in the application, and for signalling the thread that some kind of event has occurred.

## **kernel**

See **microkernel**.

## **level-sensitive**

One of two ways in which a **PIC** (Programmable Interrupt Controller) can be programmed to respond to interrupts. If the PIC is operating in level-sensitive mode, the IRQ is considered active whenever the corresponding hardware line is active. Contrast **edge-sensitive**.

## **linearly mapped**

A term indicating that a certain memory component is entirely addressable by the processor. Contrast **bank-switched**.

## **message**

A parcel of bytes passed from one process to another. The OS attaches no special meaning to the content of a message — the data in a message has meaning for the sender of the message and for its receiver, but for no one else.

Message passing not only allows processes to pass data to each other, but also provides a means of synchronizing the execution of several processes. As they send, receive, and reply to messages, processes

undergo various “changes of state” that affect when, and for how long, they may run.

### **microkernel**

A part of the operating system that provides the minimal services used by a team of optional cooperating processes, which in turn provide the higher-level OS functionality. The microkernel itself lacks filesystems and many other services normally expected of an OS; those services are provided by optional processes.

### **mount structure**

An optional, well-defined data structure (of type `iofunc_mount_t`) within an `iofunc_*`(*)* structure, which contains information used on a per-mountpoint basis (generally used only for filesystem resource managers). See also **attributes structure** and **OCB**.

### **mountpoint**

The location in the pathname space where a resource manager has “registered” itself. For example, the serial port resource manager registers mountpoints for each serial device (`/dev/ser1`, `/dev/ser2`, etc.), and a CD-ROM filesystem may register a single mountpoint of `/cdrom`.

### **mutex**

Mutual exclusion lock, a simple synchronization service used to ensure exclusive access to data shared between threads. It is typically acquired (`pthread_mutex_lock()`) and released (`pthread_mutex_unlock()`) around the code that accesses the shared data (usually a **critical section**). See also **critical section**.

### **name resolution**

In a QNX Neutrino network, the process by which the **Qnet** network manager converts an **FQNN** to a list of destination addresses that the transport layer knows how to get to.

### **name resolver**

Program code that attempts to convert an **FQNN** to a destination address.

### **NDP**

Node Discovery Protocol — proprietary QNX Software Systems protocol for broadcasting name resolution requests on a QNX Neutrino LAN.

### **network directory**

A directory in the pathname space that's implemented by the **Qnet** network manager.

### **Neutrino**

Name of an OS developed by QNX Software Systems.

### **NFS**

Network FileSystem — a TCP/IP application that lets you graft remote filesystems (or portions of them) onto your local namespace. Directories on the remote systems appear as part of your local filesystem and all the utilities you use for listing and managing files (e.g. **ls**, **cp**, **mv**) operate on the remote files exactly as they do on your local files.

### **NMI**

Nonmaskable Interrupt — an interrupt that can't be masked by the processor. We don't recommend using an NMI!

### **Node Discovery Protocol**

See **NDP**.

### **node domain**

A character string that the **Qnet** network manager tacks onto the nodename to form an **FQNN**.

### **nodename**

A unique name consisting of a character string that identifies a node on a network.

### **nonbootable**

A nonbootable OS image is usually provided for larger embedded systems or for small embedded systems where a separate, configuration-dependent setup may be required. Think of it as a second “filesystem” that has some additional files on it. Since it’s nonbootable, it typically won’t contain the OS, startup file, etc. Contrast **bootable**.

### **OCB**

Open Control Block (or Open Context Block) — a block of data established by a resource manager during its handling of the client’s *open()* function. This context block is bound by the resource manager to this particular request, and is then automatically passed to all subsequent I/O functions generated by the client on the file descriptor returned by the client’s *open()*.

### **package filesystem**

A virtual filesystem manager that presents a customized view of a set of files and directories to a client. The “real” files are present on some medium; the package filesystem presents a virtual view of selected files to the client.

### **pathname prefix**

See **mountpoint**.

### **pathname space mapping**

The process whereby the Process Manager maintains an association between resource managers and entries in the pathname space.

## **persistent**

When applied to storage media, the ability for the medium to retain information across a power-cycle. For example, a hard disk is a persistent storage medium, whereas a ramdisk is not, because the data is lost when power is lost.

## **Photon microGUI**

The proprietary graphical user interface built by QNX Software Systems.

## **PIC**

Programmable Interrupt Controller — hardware component that handles IRQs. See also **edge-sensitive**, **level-sensitive**, and **ISR**.

## **PID**

**Process ID**. Also often *pid* (e.g. as an argument in a function call).

## **POSIX**

An IEEE/ISO standard. The term is an acronym (of sorts) for Portable Operating System Interface — the “X” alludes to “UNIX”, on which the interface is based.

## **POSIX layer calls**

Convenient set of library calls for writing resource managers. The POSIX layer calls can handle even more of the common-case messages and functions than the **base layer calls**. These calls are identified by the *iofunc\_\**(*)* prefix. In order to use these (and we strongly recommend that you do), you must also use the well-defined POSIX-layer **attributes** (*iofunc\_attr\_t*), **OCB** (*iofunc\_ocb\_t*), and (optionally) **mount** (*iofunc\_mount\_t*) structures.

## **preemption**

The act of suspending the execution of one thread and starting (or resuming) another. The suspended thread is said to have been “preempted” by the new thread. Whenever a lower-priority thread is

actively consuming the CPU, and a higher-priority thread becomes READY, the lower-priority thread is immediately preempted by the higher-priority thread.

### **prefix tree**

The internal representation used by the Process Manager to store the pathname table.

### **priority inheritance**

The characteristic of a thread that causes its priority to be raised or lowered to that of the thread that sent it a message. Also used with mutexes. Priority inheritance is a method used to prevent **priority inversion**.

### **priority inversion**

A condition that can occur when a low-priority thread consumes CPU at a higher priority than it should. This can be caused by not supporting priority inheritance, such that when the lower-priority thread sends a message to a higher-priority thread, the higher-priority thread consumes CPU *on behalf of* the lower-priority thread. This is solved by having the higher-priority thread inherit the priority of the thread on whose behalf it's working.

### **process**

A nonschedulable entity, which defines the address space and a few data areas. A process must have at least one **thread** running in it — this thread is then called the first thread.

### **process group**

A collection of processes that permits the signalling of related processes. Each process in the system is a member of a process group identified by a process group ID. A newly created process joins the process group of its creator.



### **process group ID**

The unique identifier representing a process group during its lifetime. A process group ID is a positive integer. The system may reuse a process group ID after the process group dies.

### **process group leader**

A process whose ID is the same as its process group ID.

### **process ID (PID)**

The unique identifier representing a process. A PID is a positive integer. The system may reuse a process ID after the process dies, provided no existing process group has the same ID. Only the Process Manager can have a process ID of 1.

### **pty**

Pseudo-TTY — a character-based device that has two “ends”: a master end and a slave end. Data written to the master end shows up on the slave end, and vice versa. These devices are typically used to interface between a program that expects a character device and another program that wishes to use that device (e.g. the shell and the **telnet** daemon process, used for logging in to a system over the Internet).

### **pulses**

In addition to the synchronous Send/Receive/Reply services, QNX Neutrino also supports fixed-size, nonblocking messages known as pulses. These carry a small payload (four bytes of data plus a single byte code). A pulse is also one form of **event** that can be returned from an ISR or a timer. See *MsgDeliverEvent()* for more information.

### **Qnet**

The native network manager in QNX Neutrino.

## QoS

Quality of Service — a policy (e.g. **loadbalance**) used to connect nodes in a network in order to ensure highly dependable transmission. QoS is an issue that often arises in high-availability (**HA**) networks as well as realtime control systems.

## RAM

Random Access Memory — a memory technology characterized by the ability to read and write any location in the device without limitation. Contrast **flash** and **EPROM**.

## raw mode

In raw input mode, the character device library performs no editing on received characters. This reduces the processing done on each character to a minimum and provides the highest performance interface for reading data. Also, raw mode is used with devices that typically generate binary data — you don't want any translations of the raw binary stream between the device and the application. Contrast **canonical mode**.

## replenishment

In **sporadic** scheduling, the period of time during which a thread is allowed to consume its execution **budget**.

## reset vector

The address at which the processor begins executing instructions after the processor's reset line has been activated. On the x86, for example, this is the address 0xFFFFFFFF.

## resource manager

A user-level server program that accepts messages from other programs and, optionally, communicates with hardware. QNX Neutrino resource managers are responsible for presenting an interface to various types of devices, whether actual (e.g. serial ports, parallel ports, network cards, disk drives) or virtual (e.g. `/dev/null`, a network filesystem, and pseudo-ttys).

In other operating systems, this functionality is traditionally associated with **device drivers**. But unlike device drivers, QNX Neutrino resource managers don't require any special arrangements with the kernel. In fact, a resource manager looks just like any other user-level program. See also **device driver**.

## **RMA**

Rate Monotonic Analysis — a set of methods used to specify, analyze, and predict the timing behavior of realtime systems.

## **round robin**

Scheduling algorithm whereby a thread is given a certain period of time to run. Should the thread consume CPU for the entire period of its timeslice, the thread will be placed at the end of the ready queue for its priority, and the next available thread will be made READY. If a thread is the only thread READY at its priority level, it will be able to consume CPU again immediately. See also **adaptive**, **FIFO**, and **sporadic**.

## **runtime loading**

The process whereby a program decides *while it's actually running* that it wishes to load a particular function from a library. Contrast **static linking**.

## **scheduling latency**

The amount of time that elapses between the point when one thread makes another thread READY and when the other thread actually gets some CPU time. Note that this latency is almost always at the control of the system designer.

Also designated as " $T_{s1}$ ". Contrast **interrupt latency**.

## **session**

A collection of process groups established for job control purposes. Each process group is a member of a session. A process belongs to the session that its process group belongs to. A newly created process

joins the session of its creator. A process can alter its session membership via *setsid()*. A session can contain multiple process groups.

### **session leader**

A process whose death causes all processes within its process group to receive a SIGHUP signal.

### **software interrupts**

Similar to a hardware interrupt (see **interrupt**), except that the source of the interrupt is software.

### **sporadic**

Scheduling algorithm whereby a thread's priority can oscillate dynamically between a "foreground" or normal priority and a "background" or low priority. A thread is given an execution **budget** of time to be consumed within a certain **replenishment** period. See also **adaptive**, **FIFO**, and **round robin**.

### **startup code**

The software component that gains control after the IPL code has performed the minimum necessary amount of initialization. After gathering information about the system, the startup code transfers control to the OS.

### **static bootfile**

An image created at one time and then transmitted whenever a node boots. Contrast **dynamic bootfile**.

### **static linking**

The process whereby you combine your modules with the modules from the library to form a single executable that's entirely self-contained. The word "static" implies that it's not going to change — *all* the required modules are already combined into one.

### **system page area**

An area in the kernel that is filled by the startup code and contains information about the system (number of bytes of memory, location of serial ports, etc.) This is also called the SYSPAGE area.

### **thread**

The schedulable entity under QNX Neutrino. A thread is a flow of execution; it exists within the context of a **process**.

### **timer**

A kernel object used in conjunction with time-based functions. A timer is created via *timer\_create()* and armed via *timer\_settime()*. A timer can then deliver an **event**, either periodically or on a one-shot basis.

### **timeslice**

A period of time assigned to a **round-robin** or **adaptive** scheduled thread. This period of time is small (on the order of tens of milliseconds); the actual value shouldn't be relied upon by any program (it's considered bad design).



## Index

---

### !

(MQ\_PRIO\_MAX-1) 1651  
.rhosts 2518  
/dev/name/global 1768  
/dev/name/local 1768  
/dev/zero 1586  
/etc/autoconnect 1798  
/etc/hosts 464, 800, 801,  
804-807, 811, 812, 816,  
961  
/etc/hosts.equiv 2518  
/etc/networks 836, 838, 840,  
1805, 2703  
/etc/protocols 490, 862, 864,  
866, 1994, 2711  
/etc/resolv.conf 2376, 2378,  
2381, 2384, 2387, 2390  
/etc/services 492, 891, 893,  
895, 2659, 2725  
\_\_cabsargs, \_\_cabsfargs 222  
\_\_res\_state 2374  
\_client\_info 332  
\_clockadjust 281  
\_clockperiod 303  
\_cred\_info 332

\_CS\_HOSTNAME 817, 2684  
\_CS\_TIMEZONE 3342  
\_fdinfo 1172, 1177  
\_intrspin\_t 1083, 1095, 1100  
\_io\_connect 1111  
\_io\_connect\_ftime\_reply  
1118  
\_io\_connect\_link\_reply  
1120  
\_itimer 3295  
\_msg\_info 1658, 1672  
\_NTO\_SIDE\_CHANNEL 327  
\_pccard\_info 1888  
\_pulse 2238  
\_sched\_info 2574  
\_server\_info 341  
\_thread\_attr 3239  
\_thread\_local\_storage 3242  
\_timer\_info 3291

### 1

8-bit characters, reading 1601  
8086 mode, virtual 1107

**A**

- abort()* 113
- abs()* 115
- absolute values
  - complex number 222
  - floating point 574
  - integer 115
  - long integer 1421
- accept()* 117, 3189, 3573
- access()* 120
- ACCOUNTING 3394
- acos()*, *acosf()* 123
- acosh()*, *acoshf()* 125
- ACTION** 963
- addresses
  - hosts
    - strings, converting
      - to/from 1036, 1039
  - IP
    - strings, converting
      - to/from 1015, 1017, 1032, 1034, 1036, 1039
  - IPv6 1051
  - link-local 1052
  - local network
    - IP addresses, converting
      - to/from 1021
    - IP addresses, extracting
      - from 1019
  - scoped 1052
  - site-local 1052
  - sockets 127, 693, 748, 753
- addrinfo** 127
  - errors 748
  - freeing 693
  - getting 753
- advisory locks *See* files, locking
- AF\_INET 800, 803, 806, 836, 961, 983, 1036, 1039, 1313, 1805, 3189, 3348
- AF\_INET6 806, 985, 1036, 1039, 1338
- AF\_LOCAL 3377
- AF\_UNSPEC 2481
- AH (Authentication Header) 1320
- AIMS (Auto Incrementing Mass Storage) 1880
- aio\_cancel()* 129
- aio\_cb** 1443
- aio\_error()* 131
- aio\_fsync()* 133
- aio\_read()* 135
- aio\_return()* 136
- aio\_suspend()* 138
- aio\_write()* 140
- alarm()* 141
- alarms, scheduling 141, 3345
- aligned memory, allocating 1530, 1961
- alloca()* 144
- alphabetic, testing a character
  - for 1350, 1381
- alphanumeric, testing a character
  - for 1348, 1379
- alphasort()* 147
- \_amblsiz* 149, 1500, 2520, 2875
- ANSI classification 103
- arccosines 123
- architecture, instruction set 318
- arcsines 154
- arctangents 161, 163
- \_argc* 150



- argument lists,
    - variable-length 3400, 3406, 3408, 3410
    - coercion 3401
  - arguments to *main()* 150, 151, 198
    - parsing 842
  - \_argv* 151
  - arrays
    - allocating 224, 2529
    - quick-sorting 2277
  - ASCII, testing a character for 1352
  - asctime()*, *asctime\_r()* 152
  - asin()*, *asinf()* 154
  - asinh()*, *asinhf()* 156
  - assert()* 158
  - asynchronous I/O
    - canceling 129
    - error status, getting 131
    - file, synchronizing 133
    - reading 135
    - return status, getting 136
    - waiting for completion 138
    - writing 140
  - asynchronous SNMP
    - transactions 2892
  - atan()*, *atanf()* 161
  - atan2()*, *atan2f()* 163
  - atanh()*, *atanhf()* 165
  - atexit()* 167
  - atof()* 170
  - atoh()* 172
  - atoi()* 174
  - atol()*, *atoll()* 176
  - atomic operations
    - addition 178, 180
    - bits
      - clearing 182, 184
      - setting 186, 188
      - toggling 194, 196
      - subtraction 190, 192
  - atomic\_add()* 178
  - atomic\_add\_value()* 180
  - atomic\_clr()* 182
  - atomic\_clr\_value()* 184
  - atomic\_set()* 186
  - atomic\_set\_value()* 188
  - atomic\_sub()* 190
  - atomic\_sub\_value()* 192
  - atomic\_toggle()* 194
  - atomic\_toggle\_value()* 196
  - Authentication Header (AH) 1320
  - authenticator()* 2898
  - Auto Incrementing Mass Storage (AIMS) 1880
  - \_auxv* 198
- ## B
- background processes 360, 1978
    - termination, notification of 1980
  - barriers
    - attributes 2043
    - destroying 2047
    - initializing 2051
    - process-shared 2049, 2053
    - destroying 2041
    - initializing 2043
  - BARRIER\_SERIAL\_THREAD 2045
  - basename()* 199
  - bcmp()* 202
  - bcopy()* 204

- \_\_BEGIN\_DECLS 111
- Bessel functions
  - first kind 1408, 1410, 1412
  - second kind 3563, 3565, 3567
- big endian
  - \_\_BIGENDIAN\_\_ manifest 111
  - little endian, converting
    - to/from 477, 479, 481, 483, 485, 487, 3093
  - messages 1550
  - native format, converting
    - to/from 465, 467, 469
  - ports
    - reading from 1005, 1009
    - writing to 1858, 1862
  - unaligned values
    - accessing safely 3364, 3366, 3368
    - writing safely 3358, 3360, 3362
- \_\_BIGENDIAN\_\_ 111
- binary search 215
- bind()* 206, 3189, 3573
- bindresvport()* 209
- BIOS (PCI), determining if present 1922
- bits
  - atomic operations
    - clearing 182, 184
    - setting 186, 188
    - toggling 194, 196
  - set, finding first 608
- block buffering, setting for stream I/O 2662
- block special devices
  - reading 2315
  - writing 3555
- blocks
  - allocating 2875
  - reading from a file 2315
  - system message log, writing to 2867
  - writing to a file 3555
- blocksize, filesystem 714, 3001
- booting 3150
  - time since 3154
- BOOT\_TIME 929, 3393
- break condition, asserting 3192
- break pointer
  - advancing 2520
  - increment 149
- brk()* 211
- BRKINT 3211
- bsearch()* 215
- BSS data 459, 462
- \_btext* 217
- btowc()* 218
- buffers
  - canonical input 673, 1866
  - locking 1580, 1758
  - raw input 673, 1866
  - stream I/O 2660, 2742
    - block 2662
    - flushing 3173
    - line 2696
- BUFSIZ 2660
- BULK\_REQ\_MSG 2887
- busy-waiting 1784, 1786, 1789, 1791, 1793
- bytes
  - comparing 202
  - copying 204, 1532, 1538
  - overlapping objects 1544

multibyte character, number of  
     bytes in 1506, 1509  
 reading 1965, 2306  
 reordering 969, 971, 1829,  
     1831  
 setting 1546  
 writing 2272, 3550  
 zeroing 220  
*bzero()* 220

## C

C++ programs  
     end of C code 111  
     start of C code 111  
*cabs()*, *cabsf()* 222  
 calendar times  
     current 3257  
     local times, converting  
         to/from 1454, 1456, 1577  
     tm 3313  
*callback()* 2899  
*calloc()* 224  
 cancellation  
     cleanup handlers 2057, 2059  
     points 3241  
         creating 2234  
     state 2196, 3235  
     type 2198, 3235  
 canonical input buffer 673, 1866  
 C\_ANY 2377, 2380, 2383, 2386  
 Card Information Structure (CIS),  
     reading 1894  
 CardBus cards 1932  
*cbrt()*, *cbrtf()* 226

C\_CHAOS 2377, 2380, 2383, 2386  
*ceil()*, *ceilf()* 228  
*cfgetispeed()* 232  
*cfgetospeed()* 234  
 CFGFILE\_APPEND 237  
 CFGFILE\_CREAT 237  
 CFGFILE\_EXCL 237  
 CFGFILE\_RDONLY 237  
 CFGFILE\_RDWR 237  
 CFGFILE\_TRUNC 238  
 CFGFILE\_WRONLY 237  
*cfgopen()* 236  
*cfmakeraw()* 230  
*cfree()* 240  
*cfsetispeed()* 242  
*cfsetospeed()* 245  
*ChannelCreate()*,  
     *ChannelCreate\_r()* 248  
*ChannelDestroy()*,  
     *ChannelDestroy\_r()* 255  
 channels  
     creating 248, 1768  
     destroying 255, 1777  
     events, delivering 1661  
     flags 250  
     ID 249  
     messages  
         receiving 1689, 1700  
         replying 1704, 1707  
         sending 1710, 1714, 1722,  
             1726, 1730, 1734, 1738,  
             1742  
     pulses  
         receiving 1694, 1697  
         sending 1719  
     side channels 327

- character device terminal drivers,
  - providing session support to 1991
- characters *See also* strings; wide characters
  - 8-bit, reading 1601
  - control
    - disabling 674, 1867
    - discarding on input 1601
  - default (signed or unsigned) 111
  - devices
    - input stream, injecting 3184
    - size 3182, 3202
  - escape 1612
  - handling 2699
  - international *See* wide characters
  - lowercase, converting to 1601, 3321
  - multibyte
    - bytes, counting 1506, 1509
    - wide characters, conversion object 1514
    - wide characters, converting to/from 1511, 1516, 1518, 1521, 3474, 3502, 3518, 3528
  - number waiting to be read 3187
  - searching for 1534, 3015, 3058
  - sets, searching for 3025, 3056, 3068
  - special 1612
  - stdin*, reading from 611, 764, 766
  - stdout*, writing to 680, 2251, 2253
  - streams
    - pushing back 3373
    - reading from 609, 615, 760, 762
    - writing to 678, 2247, 2249
  - testing for
    - alphabetic 1350
    - alphanumeric 1348
    - ASCII 1352
    - control character 1356
    - decimal digit 1358
    - hexadecimal digit 1403
    - lowercase 1366
    - printable 1362, 1370
    - punctuation 1372
    - uppercase 1377
    - whitespace 1374
  - uppercase, converting to 3323, 3329
  - wide characters, converting to/from 218, 3526
- CHAR\_MAX 1451
- \_\_CHAR\_SIGNED\_\_ 111
- \_\_CHAR\_UNSIGNED\_\_ 111
- chdir()* 258
- chmod()* 261, 1135
  - resource managers, implementing in 1145, 1148
- chown()* 265, 1135
  - resource managers, implementing in 1150, 1153
  - restricting use of 674, 1867
- chroot()* 268

- C\_HS 2377, 2380, 2383, 2386
- chsize()* 271
- CIDR (Classless Internet Domain Routing) 1023, 1028
- C\_IN 2377, 2380, 2383, 2386
- CIS (Card Information Structure), reading 1894
- classes
  - IP addresses 1024
  - PCI 1912
  - wide-character 3533
- Classless Internet Domain Routing (CIDR) 1023, 1028
- clearenv()* 274
- clearerr()* 277
- cli** 1085
- CLK\_TCK 3307
- CLOCAL 3212
- clock
  - adjusting 282
  - CPU time, getting ID of for a thread 2102
  - cycles 284
  - getting 306
  - ID, getting 286
  - period, getting and setting 304
  - resolution, getting 288
  - setting 297, 306
  - ticks 3136
    - per seconds 3154
    - time, getting 290
- clock\_t** 3136, 3306
- clock()* 279
- ClockAdjust()*, *ClockAdjust\_r()* 282
- ClockCycles()* 284
- clock\_getcpuclockid()* 286
- clock\_getres()* 288
- clock\_gettime()* 290
- ClockId()*, *ClockId\_r()* 300
- CLOCK\_MONOTONIC 293, 3274, 3281, 3284, 3298
- clock\_nanosleep()* 294
- ClockPeriod()*, *ClockPeriod\_r()* 304
- CLOCK\_REALTIME 281, 293, 303, 306, 2130, 2173, 2176, 2641, 3259, 3274, 3281, 3284, 3298
- clock\_settime()* 297
- CLOCK\_SOFTTIME 293, 3259, 3274, 3281, 3284, 3298
- CLOCKS\_PER\_SEC 279
- ClockTime()*, *ClockTime\_r()* 306
- close()* 309
  - resource managers, implementing in 1157, 1160
- closedir()* 311, 1843
- closelog()* 314
- \_cmdfd()* 315
- \_cmdname()* 316
- cmsghdr** 3378
- code, portability 103
- collating sequence, setting 2699
- COLUMNS** 274
- commands
  - executing 3158
    - on a remote host 2295, 2488
  - options, parsing 842, 918
- communications line
  - break condition, asserting 3192
  - disconnecting 3167
- comparison

- bytes 202, 1536, 1542
  - strings
    - case-insensitive 1542, 3010, 3019, 3037
    - case-sensitive 1536, 3017
    - locale's collating sequence, using 3021
    - wide-character 3480, 3482, 3494, 3537
  - substrings
    - case-insensitive 3043, 3052
    - case-sensitive 3048
  - compiling with optimization 111
  - complementary error function 502
  - complex numbers, absolute value of 222
  - computer, rebooting 3150
  - condition variables
    - attributes
      - clock 2079, 2085
      - destroying 2077
      - initializing 2083
      - process-shared 2081, 2087
    - blocking on 2070, 2074
    - destroying 2064, 3114
    - initializing 2066, 3133
    - synchronization objects 3104, 3107
    - unblocking
      - all threads 2062
      - highest priority thread 2068
    - waiting on 2074
      - timed 2070
  - configuration files, opening 236, 577
  - configuration strings
    - \_CS\_HOSTNAME** 817, 2684
    - \_CS\_TIMEZONE** 3342
      - getting and setting 318
  - configuration values, getting 673, 1249, 1866, 3136
  - confstr()* 318
  - connect functions (resource managers)
    - default values, setting 1179
    - open, default 1241
  - connect()* 323, 3189, 3572
  - ConnectAttach()*, *ConnectAttach\_r()* 327
  - ConnectClientInfo()*, *ConnectClientInfo\_r()* 331
  - ConnectDetach()*, *ConnectDetach\_r()* 251, 335
  - ConnectFlags()*, *ConnectFlags\_r()* 337
- connections
    - client, information about 331
    - detaching 251, 335
    - dispatch interface, creating 1555
    - flags, modifying 337
    - ID 327
    - server, information about 340
    - shutting down
      - full-duplex 2762
    - sockets
      - accepting on 117, 2281
      - initiating on 323, 1798, 2298
      - listening on 1447, 2475

- ConnectServerInfo()*,
    - ConnectServerInfo\_r()* 340
  - console I/O
    - stderr* 69, 3005
    - stdin* 69, 3006
    - stdout* 69, 3007
  - const** 102
  - contention scope 2011, 2033
  - control characters 3213
    - disabling 674, 1867
    - discarding on input 1601
    - testing a character for 1356, 1383
  - controlling terminals
    - making 3200
    - path name 355
  - copysign()*, *copysignf()* 343
  - core files, maximum size 2720
  - cos()*, *cosf()* 345
  - cosh()*, *coshf()* 347
  - cosines 345
    - hyperbolic 347
    - inverse hyperbolic 125
  - CREAD 3212
  - creat()*, *creat64()* 349, 3353
  - crypt()* 353
  - CS5,...,CS8 3212
  - \_CS\_ARCHITECTURE 318
  - \_CS\_DOMAIN 318
  - \_CS\_HOSTNAME 319
  - \_CS\_HW\_PROVIDER 319
  - \_CS\_HW\_SERIAL 319
  - CSIZE 3212
  - \_CS\_LIBPATH 319
  - \_CS\_MACHINE 319
  - \_CS\_PATH 319
  - \_CS\_RELEASE 319
  - \_CS\_RESOLVE 319
  - \_CS\_SRPC\_DOMAIN 319
  - \_CS\_SYSNAME 319
  - \_CS\_TIMEZONE 319
  - CSTOPB 3212
  - \_CS\_VERSION 319
  - ctermid()* 355
  - ctime()*, *ctime\_r()* 357
  - CTL\_MAXNAME 3145
  - CTL\_NET 3140, 3141
  - cube roots 226
  - current working directory 258, 768, 940
- ## D
- daemon()* 360
  - daemons
    - SNMP (Simple Network Management Protocol), configuration file for 2311
    - system 360, 1978
    - termination, notification of 1980
  - data segment
    - changing space allocated for 211
    - end of 459, 462
    - maximum size 2720
  - data server
    - applications
      - deregistering 440
      - registering 446

- variables
  - creating 438
  - deleting 435
  - flags, setting 442
  - getting 444
  - setting 448
- data streams, flow control 3170
- databases
  - blocks
    - reading 2315
    - writing 3555
  - groups
    - closing 463
    - ID, getting information
      - about 786, 788
    - membership 1061
    - name, getting information
      - about 791, 793
    - next entry, getting 783
    - rewinding 2678
  - hosts
    - closing 464
    - entries, getting 800, 804, 806, 810, 812, 815
    - errors 953, 958, 967
    - opening 2682
    - structure 961
  - network
    - closing 489
    - entries, getting 836, 838, 840
    - opening 2703
    - structure 1805
  - passwords
    - closing 491
    - encrypting 353, 2275
    - entry, getting for a user 871, 873, 876, 878
    - entry, getting next 868
    - rewinding 2713
  - protocols
    - closing 490
    - entries, getting 862, 864, 866
    - opening 2711
    - protoent** 1994
  - services
    - closing 492
    - entries, getting 891, 893, 895
    - entry structure 2659
    - opening 2725
  - shadow passwords
    - closing 493
    - entry, reading 618, 911, 915
    - entry, structure 2260
    - entry, writing 2260
    - rewinding 2732
  - system packet forwarding *See* **ROUTE**
- datagrams 2908, 2909, 3348
- date** 3342
- daylight* 362, 3342
- daylight saving time 362, 3342
- \_DCMD\_ALL** 371
- DCMD\_ALL\_GETFLAGS** 1166
- DCMD\_ALL\_GETMOUNTFLAGS** 1166
- DCMD\_ALL\_SETFLAGS** 1166
- \_DCMD\_BLK** 372
- \_DCMD\_CAM** 372
- \_DCMD\_CHR** 372
- \_DCMD\_FSYS** 372
- \_DCMD\_INPUT** 372



- .\_DCMD\_IP 372
- .\_DCMD\_MEM 372
- .\_DCMD\_MISC 372
- .\_DCMD\_MIXER 372
- .\_DCMD\_NET 372
- .\_DCMD\_PHOTON 372
- .\_DCMD\_PROC 372
- DEAD\_PROCESS 929, 3394
- DebugBreak()* 363
- debugging
  - kernel 365, 366
  - printing debugging messages  
(resolver routines) 2374
  - processes 363
  - shared objects 422
  - sockets 904
- DebugKDBreak()* 365
- DebugKDOutput()* 366
- decimal digit, testing a character
  - for 1358, 1387
- decimal-point character,
  - setting 2699
- delay()* 368
- DELAYTIMER\_MAX 3267
- dev\_t** 2511
- devctl()* 371
  - resource managers,  
implementing in 1166,  
1170
- \_DEVCTL\_DATA()* 1168
- DEVDIR\_FROM 372
- DEVDIR\_NONE 372
- DEVDIR\_TO 372
- DEVDIR\_TOFROM 372
- devices
  - block special  
reading 2315
  - writing 3555
- character
  - characters, injecting 3184
  - size 3182, 3202
- classes of 2509
- controlling 371, 1123
- controlling device,
  - making 3200
- I/O memory, mapping 1591,  
1764
- mounts, autodetecting 1618
- numbers
  - attaching 2510
  - detaching 2513
  - getting 2417
  - manipulating 2997
  - opening 1597
  - output, waiting for  
completion 3165
- PCI
  - attaching 1899
  - configuration, reading 1924,  
1926, 1928, 1930
  - configuration, writing 1934,  
1937, 1939, 1941
  - detaching 1910
  - finding 1912, 1914
  - rescanning for 1932
- reading 1602, 2318
- script, running on 1605
- writing 1612
- devp-pccard** server
  - arming 1881
  - attaching 1884
  - card insertion/removal,  
notification of 1881

- CIS (Card Information Structure), reading 1894
- detaching 1886
- locking 1891
- socket setup information 1888
- unlocking 1895
- \_DEXTRA\_FIRST()* 2325
- \_DEXTRA\_NEXT()* 2325
- \_DEXTRA\_VALID()* 2325
- D\_FLAG\_FILTER 382
- D\_FLAG\_STAT 382
- difftime()* 380
- digit, testing a character for
  - decimal 1358, 1387
  - hexadecimal 1401, 1403
- \_DIOF()* 371
- \_DION()* 371
- \_DIOT()* 371
- \_DIOTF()* 371
- DIR 2324
- direntl()* 382
- direct memory access (DMA)
  - channels, managing 2494
- directories
  - access, checking 120, 456
  - base name 199
  - closing 311
  - controlling 382
  - creating 1563, 1569
  - current working 258, 768, 940
    - for daemons 1978
  - deleting 2477
  - entries
    - duplicate, filtering 382
    - sorting 147
  - files, searching for 1871
  - hierarchy, walking 731, 1820
  - information, requesting 382
  - name 385
  - opening 1843
  - position
    - getting 3207
    - setting 2607
  - reading 2324, 2328
  - rewinding 2465
  - root 268
  - scanning 2531
- dirent** 2324, 2328
- dirname()* 385
- dispatch\_t** 396
- dispatch interface *See also*
  - resource managers
    - blocking 388
  - connections, creating 1555
  - context
    - allocating 391
    - freeing 394
  - endian, converting 1550
  - events
    - handling 401
    - last selected 2623
  - file descriptors
    - handle, attaching 2616
    - handle, detaching 2619
  - handles
    - creating 396
    - destroying 399
  - message handlers
    - attaching 1549
    - detaching 1558
  - names
    - attaching 1768
    - detaching 1777

- server connections,
  - closing 1775
  - server connections,
    - opening 1779
- path
  - attaching to 2394
  - detaching from 2413
- pulse handlers
  - attaching 2241
  - detaching 2244
- thread pool
  - attributes, changing 3216, 3229
  - creating 3218
  - destroying 3225
  - starting 3231
  - timeout, setting 404
  - unblocking 406
- dispatch\_block()* 388
- dispatch\_context\_alloc()* 391
- dispatch\_context\_free()* 394
- dispatch\_create()* 396
- dispatch\_destroy()* 399
- dispatch\_handler()* 401
- dispatch\_timeout()* 404
- dispatch\_unblock()* 406
- div\_t** 408
- div()* 408
- division
  - integer 408
  - long integer 1430
- DL\_info** 410
- dladdr()* 411
- dlclose()* 413
- DL\_DEBUG** 422
- dLError()* 415
- dlopen()* 417
- dlsym()* 424
- DMA channels, managing 2494
- dn\_comp()* 427
- dn\_expand()* 429
- domains
  - errors 953, 958, 967
  - names
    - compressing 427
    - expanding 429
    - getting 318, 771
    - resolving 2374, 2378, 2381, 2384, 2387, 2389
    - setting 2664
  - secure RPC 319
  - UNIX 3377
- dot notation (IP addresses) 1024
- dotted quad (IP addresses) 1024
- double-precision numbers
  - absolute value 222, 574
  - arccosines 123
  - arcsines 154
  - arctangents 161, 163
  - Bessel functions 1408, 1410, 1412, 3563, 3565, 3567
  - complementary error
    - function 502
  - cosines 345
  - cube roots 226
  - error function 500
  - exceptions, signal for 2797
  - exponentials 569, 571
  - exponents,
    - radix-independent 2523, 2526
  - finite, determining if 629

- fractional part of a
  - double-precision number 1615
- gamma functions 750, 1435
- hyperbolic cosines 347
- hyperbolic sines 2849
- hyperbolic tangents 3163
- hypotenuse, length of 981
- infinite, determining if 1364
- input, formatted 2533
- integral logarithms 999
- integral part of a
  - double-precision number 1615
- integral power of 2 701, 1428
- inverse hyperbolic cosines 125
- inverse hyperbolic sines 156
- inverse hyperbolic tangents 165
- logarithms 1462, 1464, 1466
- modular arithmetic 643
- next representable 1816
- normalized fractions 701
- not a number, determining
  - if 1368
- powers 1963
- precision 667
- printing 1968
- pseudo-random numbers 431, 498
- radix-independent
  - exponents 1468, 2523, 2526
- remainders 433, 2366
- residue 643
- rounding 228, 639, 2472
- sign, copying 343
- significant bits 2819
- sines 2847
- square roots 2980
- strings, converting
  - to/from 170, 3072
- tangents 3161
- times, difference between 380
- wide-character strings,
  - converting to/from 3508
- drand48()* 431
- drem()*, *dremf()* 433
- ds\_clear()* 435
- ds\_create()* 438
- ds\_deregister()* 440
- ds\_flags()* 442
- ds\_get()* 444
- DS\_PERM 437, 440, 442
- ds\_register()* 446
- ds\_set()* 448
- \_DTYPE\_LSTAT 1258, 1305, 2325
- \_DTYPE\_NONE 2325
- \_DTYPE\_STAT 2325
- dup()* 450
- dup2()* 453
- duplex connection, shutting
  - down 2762
- dynamically linked libraries
  - addresses, translating 411
  - closing 413
  - debugging 422
  - errors 415
  - opening 417
  - symbol, getting address of 424

**E**

- E2BIG 508
- EACCES 508
- eaccess()* 456
- EADDRINUSE 508
- EADDRNOTAVAIL 508
- EADV 508
- EAFNOSUPPORT 508, 2913
- EAGAIN 508
- EALREADY 508
- EBADE 508
- EBADF 508, 1881, 1892, 1896
- EBADFD 508
- EBADFSYS 508
- EBADMSG 508
- EBADR 508
- EBADRPC 508
- EBADRQC 508
- EBADSLT 508
- EBFONT 508
- EBUSY 508, 1892
- ECANCELED 508
- ECHILD 508
- ECHO 3212
- ECHOE 3212
- ECHOK 3212
- ECHONL 3213
- ECHRNG 508
- ECOMM 508
- ECONNABORTED 508
- ECONNREFUSED 508, 1447
- ECONNRESET 508
- ECTRLTERM 508
- \_edata* 459
- EDEADLK 508
- EDEADLOCK 508
- EDESTADDRREQ 508
- EDOM 508
- EDQUOT 508
- EEXIST 508
- EFAULT 508
- EFBIG 508
- EHOSTDOWN 508
- EHOSTUNREACH 508
- EIDRM 508
- EILSEQ 508
- EINPROGRESS 508
- EINTR 508
- EINVAL 508
- EIO 508
- EISCONN 508, 3572
- EISDIR 508
- EL2HLT 508
- EL2NSYNC 508
- EL3HLT 508
- EL3RST 508
- ELIBACC 508
- ELIBBAD 508
- ELIBEXEC 508
- ELIBMAX 508
- ELIBSCN 508
- ELNRNG 508
- ELOOP 508
- EMFILE 508
- EMLINK 508
- EMORE 508
- EMPTY 3393
- EMSGSIZE 508
- EMULTIHOP 508
- ENAMETOOLONG 508
- Encapsulated Security Payload  
(ESP) 1320
- encrypt()* 460

- encryption
  - key, setting 2694
  - passwords 353, 2275
  - strings 460
- \_end* 462
- end-of-file
  - files 496
  - streams 602
    - clearing 277
- \_END\_DECLS* 111
- endgrent()* 463
- endhostent()* 464
- endian
  - big
    - \_BIGENDIAN\_* manifest 111
    - little endian, converting
      - to/from 477, 479, 481, 483, 485, 487, 3093
    - messages 1550
    - native format, converting
      - to/from 465, 467, 469
    - unaligned values, accessing
      - safely 3364, 3366, 3368
    - unaligned values, writing
      - safely 3358, 3360, 3362
  - little
    - big endian, converting
      - to/from 477, 479, 481, 483, 485, 487, 3093
    - \_LITTLEENDIAN\_* manifest 111
    - messages 1550
    - native format, converting
      - to/from 471, 473, 475
    - unaligned values, accessing
      - safely 3364, 3366, 3368
    - unaligned values, writing
      - safely 3358, 3360, 3362
- ports
  - reading from 1005, 1009
  - writing to 1858, 1862
- ENDIAN\_BE16()* 465
- ENDIAN\_BE32()* 467
- ENDIAN\_BE64()* 469
- ENDIAN\_LE16()* 471
- ENDIAN\_LE32()* 473
- ENDIAN\_LE64()* 475
- ENDIAN\_RET16()* 477
- ENDIAN\_RET32()* 479
- ENDIAN\_RET64()* 481
- ENDIAN\_SWAP16()* 483
- ENDIAN\_SWAP32()* 485
- ENDIAN\_SWAP64()* 487
- endnetent()* 489
- endprotoent()* 490
- endpwent()* 491
- endservent()* 492
- endspent()* 493
- endutent()* 494
- ENETDOWN 508
- ENETRESET 508
- ENETUNREACH 508
- ENFILE 508
- ENOANO 508
- ENOBUFS 508, 2482
- ENOCCSI 508
- ENODATA 508
- ENODEV 508, 1890, 1892, 1896
- ENOENT 508
- ENOEXEC 508
- ENOLCK 508
- ENOLIC 508
- ENOLINK 508

- ENOMEM 508
- ENOMSG 508
- ENONDP 508
- ENONET 508
- ENOPKG 508
- ENOPROTOOPT 508
- ENOREMOTE 508
- ENOSPC 508
- ENOSR 508
- ENOSTR 508
- ENOSYS 508
- ENOTBLK 508
- ENOTCONN 508
- ENOTDIR 508
- ENOTEMPTY 508
- ENOTSOCK 508
- ENOTSUP 508
- ENOTTY 508
- ENOTUNIQ 508
- ENTER 963
- ENTRY** 963
- env** 537, 560, 2669, 2934, 2944, 2961, 2970
- environ* 495, 1496, 2670
- environment
  - restoring 1473, 2787
  - saving 2691, 2832
- environment variables
  - COLUMNS** 274
  - defining 495, 537, 559, 2255, 2669, 2927, 2939, 2948, 2954, 2966, 2974
  - deleting 274, 2255, 2669, 2927, 2939, 2948, 2954, 2966, 2974, 3383
  - DL\_DEBUG** 422
  - files, searching for 2602
  - getting 777
  - HOSTALIASES** 807
  - HOSTNAME** 817, 2684
  - INCLUDE** 2602
  - LD\_LIBRARY\_PATH** 422
  - LIB** 2602
  - LINES** 274
  - LOCALDOMAIN** 2374, 2378, 2382, 2384, 2387, 2390
  - PATH** 57, 274, 536, 559, 2602, 2943, 2947, 2952, 2969, 2973
  - pointer to 495
  - SHELL** 274, 1955, 3158
  - SNMPCONFIGFILE** 2313
  - TERM** 274
  - TERMINFO** 274
  - TZ** 274, 3342
- ENXIO 508
- EOF 3417
- eof()* 496
- EOK 508
- EOPNOTSUPP 508, 2913
- EOVERFLOW 508
- EPERM 508
- EPFNOSUPPORT 508
- EPIPE 508
- EPROCUNAVAIL 508
- EPROGMISMATCH 508
- EPROGUNAVAIL 508
- EPROTO 508
- EPROTONOSUPPORT 508, 2913
- EPROTOTYPE 508
- erand48()* 498
- ERANGE 508
- EREMCHG 508

- EREMOTE 508
- ERESTART 508, 1669
- erf()*, *erff()* 500
- erfc()* 502
- erfcf()* 502
- EROFS 508
- ERPCMISMATCH 508
- err()*, *errx()* 504
- errno* 507
  - threads 507, 3242
- error function 500
  - complementary 502
- errors *See also* signals
  - command-line options, printing
    - for 843
  - end-of-file 602
  - errno* global variable 507
  - hosts 953, 958, 967
  - message-passing 1669
  - messages
    - for an error code 1945, 3029
    - formatted 504, 3414
  - regular expressions 2359
  - resolver 953, 958, 967
  - signals, raising 2283
  - socket address
    - information 748
  - stderr* 69, 504, 1945, 3005, 3414
  - stream I/O
    - clearing 277
    - testing for 604
- escape characters 1612
- ESHUTDOWN 508
- ESOCKTNOSUPPORT 508
- ESP (Encapsulated Security Payload) 1320
- ESPIPE 508
- ESRCH 508
- ESRMNT 508
- ESRVRFAULT 508
- ESTALE 508
- ESTRPIPE 508
- .etext* 515
- ETIME 508
- ETIMEDOUT 508
- ETOOMANYREFS 508
- ETXTBSY 508
- EUNATCH 508
- EUSERS 508
- events
  - blocking while waiting for 388
  - channels, delivering
    - through 1661
  - checking validity of 1746
  - handling 401
  - interrupts
    - attaching 1077
    - detaching 1083
  - last selected 2623
  - sigevent** 2778
  - system
    - notification of 1981
    - triggering 1985
  - tracing 3331
  - unblocking 406
- EWOULDBLOCK 508, 907
- exceptional conditions, file
  - descriptors with 2493, 2610
- exceptions, floating-point
  - mask 661
  - registers 664
- exclusive locks 634



- EXDEV 508  
*exec\** family of functions 49, 57,  
     1843  
*execl()* 516  
*execle()* 522  
*execlp()* 529  
*execlpe()* 536  
 executable files  
     base name 1993  
     file descriptor 315  
     full path 316  
     mounted filesystem, preventing  
         from loading on 1617  
*execv()* 540  
*execve()* 546  
*execvp()* 552  
*execvpe()* 559  
 EXFULL 508  
*\_exit()* 563  
*exit()* 566  
*exp(), expf()* 569  
*expm1(), expm1f()* 571  
 exponentials, floating point 569,  
     571, 701, 1428  
 exponents,  
     radix-independent 1468,  
     2523, 2526  
**export** 537, 560, 2669, 2934,  
     2944, 2961, 2970
- F**
- fabs(), fabsf()* 574  
 F\_ALLOCS 588, 1268  
*fcfgopen()* 577  
*fchmod()* 578  
*fchown()* 581  
 FCHR\_MAX 2723  
*fclose()* 584  
*fcloseall()* 586  
*fcntl()* 588  
     F\_DUPFD 450, 453  
     resource managers,  
         implementing in 1190  
 FD *See* file descriptors  
*fdatasync()* 597  
 FD\_CLOEXEC 590, 2754, 2925,  
     2952  
*FD\_CLR()* 2611  
*fdinfo()*  
     resource managers,  
         implementing in 1173,  
         1175  
 \_FDINFO\_FLAG\_LOCALPATH 1125,  
     1176  
*FD\_ISSET()* 2611  
*fdopen()* 599  
*FD\_SET()* 2611  
 FD\_SETSIZE 2611  
 F\_DUPFD 589  
*FD\_ZERO()* 2611  
*feof()* 602  
*ferror()* 604  
*fflush()* 606  
 F\_FREESP 589, 1268  
*ffs()* 608  
*fgetc()* 609  
*fgetchar()* 611  
 F\_GETFD 589  
 F\_GETFL 589  
 F\_GETLK 589  
*fgetpos()* 613

- fgets()* 615, 889
- fgetspent()* 618
- fgetwc()* 621
- fgetws()* 623
- FIFO scheduling 2031
- FIFOs
  - creating 1566, 1569
  - reading from 1965, 2306, 2334
  - unnamed (pipes)
    - closing 1943
    - creating 1947
    - opening 1955
- \_\_FILE\_\_ 158
- file descriptors
  - connection IDs as 327
  - creating 1835, 2918
  - duplicating 450, 453, 589, 1846, 2923
  - exceptional conditions 2493, 2610
  - full path 3337, 3339
  - handle
    - attaching 2616
    - detaching 2619
  - maximum number of 2720
  - properties 1360
  - ready for reading or writing 2493, 2610
  - selecting 2493, 2610
  - sets of, manipulating 2611
  - stderr* 3005
  - stdin* 3006
  - stdout* 3007
  - streams, associating with 599, 626
  - table size, getting 773
  - terminal, testing for association
    - with 1354
- fileno()* 69, 626
- files
  - access times 3390
  - disabling logging of on mounted filesystems 1617
  - resource managers, implementing in 1283, 1295, 1298
  - access, checking 120, 456
  - base name 199
  - closing 309
  - configurable limits 673, 1249, 1866
  - configuration, opening 236, 577
  - controlling 588
  - core, maximum size 2720
  - creating 650, 1835, 2918
    - low-level 349
    - not allowing on mounted filesystems 1617
    - resource manager 1113, 1236
  - deleting 2368, 3380
  - device parameters, manipulating 1123
  - directory name 385
  - executable
    - base name 1993
    - file descriptor 315
    - full path 316
  - extending 271, 588
  - flushing 606
    - all 641

- information, getting 710, 1485, 2993
- input, formatted 703, 746, 3421, 3426
- link count, maximum 673, 1866
- linking to 631, 1438, 1874, 1876, 2331, 3099
- locking 589, 634, 1459
  - by a thread 637
  - nonblocking 729
- modification times 736, 3387, 3390
  - resource managers, implementing in 1283, 1295, 1298
- names
  - full path 3337, 3339
  - matching 647
  - maximum length 673, 1866
- opening 349, 650, 1835, 2918
  - resource manager 1236
- output, formatted 676, 741, 3418, 3424
- ownership, changing 265, 581, 1423
- pathnames matching a
  - pattern 943, 947
- `_PATH_UTMP` 494
- `_PATH_UTMP` 927, 930, 932, 2264, 2740, 3395
- permissions
  - changing 261, 578
  - daemons 1978
  - on creation 3353
  - restricting the changing of 674, 1867
- position
  - getting 613
  - setting 706, 708, 1482, 3204
- private access 1846
- processes, maximum files
  - per 3137
- reading 1066, 1965, 2306
  - blocks 2315
  - characters, number waiting to be read 3187
  - checking 2300
  - `iov_t` 2334
- renaming 2371
- reopening 697
- rewinding 2462
- scanning directories for 2531
- searching
  - environment variables 2602
  - list of directories 1871
- seeking 706
- shared access 2918, 2923
- size, changing 271
- size, maximum 2720
- SNMP configuration 2311
- status flags 589
- status-change times
  - resource managers, implementing in 1283, 1295, 1298
- synchronizing 597, 718
- temporary
  - creating 3315
  - creating and opening 1573
  - name, generating 1575, 3209, 3318
- tree, walking 731, 1820

- truncating 271, 589, 726, 1491, 3334
- unlinking 2368, 3380
- unlocking 590, 634, 734, 1459
- writing 2272, 3550
  - blocks 3555
  - characters 678, 680
  - `iov_t` 3558
  - strings 682
  - wide characters 684
  - wide-character strings 686
- filesystems
  - free space 715
  - information, getting 714, 3001
  - mounting 1618
    - options, parsing 1620
  - read-only 1617
  - synchronizing 3102
    - requesting 1985
  - unmounting 1618, 3356
- FIND 963
- finite number, determining if 629
- `finite()`, `finitef()` 629
- first-in first-out scheduling 2031
- `fink()` 631
- floating point
  - absolute value 222, 574
  - arccosines 123
  - arcsines 154
  - arctangents 161, 163
  - Bessel functions 1408, 1410, 1412, 3563, 3565, 3567
  - complementary error
    - function 502
  - cosines 345
  - cube roots 226
  - error function 500
  - exceptions, signal for 2797
  - exponentials 569, 571
  - exponents,
    - radix-independent 2523, 2526
  - finite, determining if 629
  - fractional part 1615
  - gamma functions 750, 1435
  - hyperbolic cosines 347
  - hyperbolic sines 2849
  - hyperbolic tangents 3163
  - hypotenuse, length of 981
  - infinite, determining if 1364
  - input, formatted 2533
  - integral logarithms 999
  - integral part 1615
  - integral power of 2 701, 1428
  - inverse hyperbolic cosines 125
  - inverse hyperbolic sines 156
  - inverse hyperbolic
    - tangents 165
  - logarithms 1462, 1464, 1466
  - modular arithmetic 643
  - next representable 1816
  - normalized fractions 701
  - not a number, determining
    - if 1368
  - powers 1963
  - printing 1968
  - radix-independent
    - exponents 1468, 2523, 2526
  - remainders 433, 2366
  - residue 643
  - rounding 228, 639, 2472
  - settings
    - exception mask 661

- exception registers 664
- precision 667
- rounding 670
- sign, copying 343
- significant bits 2819
- sines 2847
- square roots 2980
- tangents 3161
- wide-character strings,
  - converting to/from 3508
- flock** 591
- flock()* 634
- flockfile()* 637
- floor()*, *floorf()* 639
- flow control 3170
- flushall()* 641
- fmod()*, *fmodf()* 643
- fnmatch()* 647
- FNM\_PATHNAME 646
- FNM\_PERIOD 646
- FNM\_QUOTE 646
- F\_OK 120, 456
- fopen()* 650
- fork()* 655, 1843
- forkpty()* 659
- fpathconf()* 673
- \_FP\_EXC\_DENORMAL 661, 664
- \_FP\_EXC\_DIVZERO 661, 664
- fp\_exception\_mask()* 661
- fp\_exception\_value()* 664
- \_FP\_EXC\_INEXACT 661, 664
- \_FP\_EXC\_INVALID 661, 664
- \_FP\_EXC\_OVERFLOW 661, 664
- \_FP\_EXC\_UNDERFLOW 661, 664
- fpos\_t** 613, 708
- \_FP\_PREC\_DOUBLE 667
- \_FP\_PREC\_DOUBLE\_EXTENDED 667
- \_FP\_PREC\_EXTENDED 667
- \_FP\_PREC\_FLOAT 667
- fp\_precision()* 667
- fprintf()* 676
- fp\_rounding()* 670
- \_FP\_ROUND\_NEAREST 670
- \_FP\_ROUND\_NEGATIVE 670
- \_FP\_ROUND\_POSITIVE 670
- \_FP\_ROUND\_ZERO 670
- fputc()* 678
- fputchar()* 680
- fputs()* 682
- fputwc()* 684
- fputws()* 686
- FQNN (Fully Qualified Node Name) 1807
- FQPN (Fully Qualified Path Name) 1806
- fractional part of a floating-point number 1615
- F\_RDLCK 590
- fread()* 688
- free memory, amount of 2999
- free space, filesystem 715, 3002
- free()* 691
- freeaddrinfo()* 693
- freeifaddrs()* 695
- freopen()* 69, 697
- frexp()*, *frexpf()* 701
- fscanf()* 703
- fseek()* 706, 2462
- F\_SETFD 590
- F\_SETFL 590
- F\_SETLK 590
- F\_SETLKW 590
- fsetpos()* 708
- fstat()*, *fstat64()* 710

*fstatvfs()*, *fstatvfs64()* 714  
*fsync()* 718  
*ftell()* 720  
*ftime()* 723  
*ftruncate()*, *ftruncate64()* 726  
*ftrylockfile()* 729  
*ftw()* 731  
FTW\_D 732, 1820  
FTW\_DNR 732, 1820  
FTW\_F 732, 1820  
FTW\_NS 732, 1820  
\_FTYPE\_ANY 1112, 1118, 1120,  
2392  
\_FTYPE\_LINK 1112, 1118, 1120,  
2392  
\_FTYPE\_MOUNT 1112, 1118, 1120,  
2392  
\_FTYPE\_QUEUE 1112, 1118,  
1120, 2392  
\_FTYPE\_PIPE 1112, 1118, 1120,  
2392  
\_FTYPE\_SEM 1112, 1118, 1121,  
2393  
\_FTYPE\_SHMEM 1112, 1118, 1121,  
2393  
\_FTYPE\_SOCKET 1112, 1118,  
1121, 2393  
\_FTYPE\_SYMLINK 1112, 1118,  
1121, 2393  
full-duplex connection, shutting  
down 2762  
Fully Qualified Node Name  
(FQNN) 1807  
Fully Qualified Path Name  
(FQPN) 1806  
function  
classification 103

safety 107  
F\_UNLCK 589, 590  
*funlockfile()* 734  
*futime()* 736  
*fwide()* 739  
*fwprintf()* 741  
*fwrite()* 743  
F\_WRLCK 590  
*fwscanf()* 746

## G

*gai\_strerror()* 748  
gamma functions 750, 1435  
*gamma()*, *gamma\_r()*, *gammaf()*,  
*gammaf\_r()* 750  
General Purpose Interface Bus  
(GPIB) 1880  
*getaddrinfo()* 753  
*getc()* 760  
*getchar()* 764  
*getchar\_unlocked()* 766  
*getc\_unlocked()* 762  
*getcwd()* 768  
*get\_device\_command()* 372  
*get\_device\_direction()* 372  
*getdomainname()* 771  
*getdtablesize()* 773  
*getegid()* 775  
*getenv()* 537, 560, 777, 2934, 2944,  
2961, 2970  
*geteuid()* 779  
*getgid()* 781  
*getgrent()* 783  
*getgrgid()* 786

*getrgid\_r()* 788  
*getrnam()* 791  
*getrnam\_r()* 793  
*getgrouplist()* 796  
*getgroups()* 798  
*gethostbyaddr()* 800  
*gethostbyaddr\_r()* 804  
*gethostbyname()*, *gethostbyname2()*  
     806  
*gethostbyname\_r()* 810  
*gethostent()* 812  
*gethostent\_r()* 815  
*gethostname()* 817  
*getifaddrs()* 819  
*GETIOVBASE()* 821  
*GETIOVLEN()* 823  
*getitimer()* 825  
*getlogin()* 827  
*getlogin\_r()* 829  
*getnetbyaddr()* 836  
*getnetbyname()* 838  
*getnetent()* 840  
*GETNEXT\_REQ\_MSG* 2887  
*getopt()* 842  
*getpass()* 848  
*getpeername()* 850  
*getpgid()* 852  
*getpgrp()* 854  
*getpid()* 856  
*getppid()* 858  
*getprio()* 860  
*getprotobyname()* 862  
*getprotobyname\_r()* 864  
*getprotoent()* 866  
*getpwent()* 868  
*getpwnam()* 871  
*getpwnam\_r()* 873  
  
*getpwuid()* 876  
*getpwuid\_r()* 878  
*GET\_REQ\_MSG* 2887  
*getrlimit()*, *getrlimit64()* 881  
*GET\_RSP\_MSG* 2887  
*getrusage()* 884  
*gets()* 889  
*getservbyname()* 891  
*getservbyport()* 893  
*getservent()* 895  
*getsid()* 897  
*getsockname()* 899, 3573  
*getsockopt()* 902  
*getspent()* 911  
*getspnam()*, *getspnam\_r()* 915  
*getsubopt()* 918  
*gettimeofday()* 923  
*getuid()* 925  
*getutent()* 927  
*getutid()* 929  
*getutline()* 932  
*getw()* 934  
*getwc()* 936  
*getwchar()* 938  
*getwd()* 940  
**glob\_t** 942  
*glob()* 943  
 global variables  
     *\_amblksiz* 149  
     *\_argc* 150  
     *\_argv* 151  
     *\_auxv* 198  
     *\_btext* 217  
     *daylight* 362  
     *\_edata* 459  
     *\_end* 462  
     *errno* 507

- \_etext* 515
- optarg* 842
- opterr* 843
- optind* 842
- optopt* 843
- \_\_progname* 1993
- stderr* 3005
- stdin* 3006
- stdout* 3007
- sys\_errlist* 507
- sys\_nerr* 508
- sys\_nsig* 2796
- \_\_syspage\_ptr* 3157
- sys\_siglist* 2796
- timezone* 3312
- tzname* 3341
- globfree()* 947
- gmtime()* 949
- gmtime\_r()* 951
- GPIB (General Purpose Interface Bus) 1880
- greater of two numbers 1504
- groups *See also* process groups
  - access list
  - getting 796
  - initializing 1061
- database
  - closing 463
  - membership 1061
  - next entry, getting 783
  - rewinding 2678
- IDs
  - effective 775, 2666, 2675, 2714
  - information about 786, 788
  - process 854
  - real 781, 2675, 2714

- saved 2675
- name, getting information
  - about 791, 793
- set-group ID 3137
- supplementary
  - IDs 798, 2680
  - maximum per process 3136
- guard area 2003, 2024
- guardian process, specifying 1987

## H

- hardware
  - information in system
    - page 973, 975, 977, 979
  - manufacturer, getting 319
  - serial number, getting 319
  - type 319, 3370
- hardware interrupts *See* interrupts
- hash table
  - creating 955
  - destroying 957
  - searching 963
- hcreate()* 955
- hdestroy()* 957
- h\_errno* 801, 807, 953, 958, 2381, 2387
- herror()* 958
- hexadecimal numbers
  - digit, testing a character
    - for 1401, 1403
  - strings, converting to/from 172
- host-byte order



- network-byte order, converting
    - to/from 969, 971, 1829, 1831
  - HOSTALIASES** 807
  - hostent** 961
  - HOSTNAME** 817, 2684
  - HOST\_NOT\_FOUND** 953, 958
  - hosts
    - addresses
      - strings, converting
        - to/from 1036, 1039
    - database
      - closing 464
      - entries, getting 800, 804, 806, 810, 812, 815
      - errors 953, 958, 967
      - opening 2682
      - structure 961
    - domain name
      - getting 318, 771
      - setting 2664
    - names 817, 2684
      - getting 319
      - valid characters 319
    - remote
      - identity, checking 2518
  - hot swapping 1932
  - hsearch()* 963
  - hstrerror()* 967
  - htonl()* 969
  - htons()* 971
  - HUPCL** 3212
  - hwi\_find\_item()* 973
  - hwi\_find\_tag()* 975
  - HWL\_NULL\_OFF** 974
  - hwi\_off2tag()* 977
  - hwi\_tag2off()* 979
  - hyperbolic functions *See also*
    - trigonometry
      - hyperbolic cosine 347
      - hyperbolic sine 2849
      - hyperbolic tangent 3163
      - inverse hyperbolic cosine 125
      - inverse hyperbolic sine 156
      - inverse hyperbolic tangent 165
    - hypot()*, *hypotf()* 981
    - hypotenuse, length of 981
- I**
- I/O**
- buffers, flushing 3173
  - configuration files,
    - opening 236
  - end-of-file, checking for 496
  - FIFOs, creating 1566, 1569
  - file descriptors
    - duplicating 450, 453, 589, 1846, 2923
    - selecting 2493, 2610
  - file-mode creation mask 3353
    - daemons 1978
  - files
    - closing 309
    - controlling 588
    - extending 588
    - information, getting 710
    - linking to 1438, 1874, 1876, 2331, 3099
    - locking 589, 1459
    - names, matching 647
    - opening 349, 1835, 2918

- reading 1965, 2306, 2315
- status flags 589
- truncating 589
- unlocking 590, 634, 1459
- writing 2272, 3550, 3555
- filesystems
  - information, getting 714, 3001
- iov\_t**
  - reading 2334
  - writing 3558
- ports, managing 2494
- privity, requesting 1069, 1078, 1083, 1095, 1097, 1100, 1102, 3246
- requests, initiating list of 1443
- I/O functions (resource managers)
  - chmod* 1145
    - default 1148
  - chown* 1150
    - default 1153
  - close*
    - default 1157, 1160
  - default values, setting 1179
  - devctl* 1166
    - default 1170
  - fdinfo* 1173
    - default 1175
  - link* 1182
  - lock*
    - default 1190
  - lseek* 1197
    - default 1200
  - mknod* 1202
  - mmap* 1205
    - default 1209
  - openfd* 1243
    - default 1247
  - pathconf* 1249
    - default 1252
  - read* 1256
    - default 1254
  - readlink* 1260
  - space* 1267
  - stat* 1271
    - default 1273
  - sync* 1280
    - default 1278
  - utime* 1295
    - default 1298
  - write*
    - default 1301
- I/O vector
  - base, getting 821
  - fields, filling 2686
  - length, getting 823
  - reading from a file 2334
  - writing to a file 3558
- ICANON 3213
- ICMP (Internet Control Message Protocol) 983
- ICMP6 (Internet Control Message Protocol v6) 985
- ICMP6\_FILTER 986
- ICMP6\_FILTER\_SETBLOCK()* 986
- ICMP6\_FILTER\_SETBLOCKALL()* 986
- ICMP6\_FILTER\_SETPASS()* 986
- ICMP6\_FILTER\_SETPASSALL()* 986
- ICMP6\_FILTER\_WILLBLOCK()* 986
- ICMP6\_FILTER\_WILLPASS()* 986
- ICRNL 3211
- IEXTEN 3213
- if\_msghdr** 2483
- if\_nameindex** 993

- ifa\_msghdr** 2483
- ifaddrs** 997
- if\_freenameindex()* 989
- if\_indexoname()* 991
- if\_nameindex()* 993
- if\_nametoindex()* 995
- IFNAMSIZ 991
- IGNBRK 3211
- IGNCR 3211
- IGNPAR 3211
- IHFLOW 3212
- ilogb(), ilogbf()* 999
- in16()* 1005
- in16s()* 1007
- in32()* 1009
- in32s()* 1011
- in6\_pktinfo** 1342
- in8()* 1001
- in8s()* 1003
- INADDR\_ANY 1315, 3189
- INADDR\_NONE 1015, 1016, 1030
- inbe16()* 1005
- inbe32()* 1009
- INCLUDE** 2602
- incoming connections, listening
  - for 1447, 2475
- index()* 1013
- INET6 (Internet protocol v6
  - family) 1051
- INET6\_ADDRSTRLEN 1036
- inet\_addr()* 1015
- INET\_ADDRSTRLEN 1036
- inet\_aton()* 1017
- inet\_lnaof()* 1019
- inet\_makeaddr()* 1021
- inet\_net\_ntop()* 1023
- inet\_netof()* 1026
- inet\_net\_pton()* 1028
- inet\_network()* 1030
- inet\_ntoa()* 1032
- inet\_ntoa\_r()* 1034
- inet\_ntop()* 1036
- inet\_pton()* 1039
- infinite number, determining
  - if 1364
- INFORM\_REQ\_MSG 2887
- inheritance** 2925, 2952
- initgroups()* 1061
- INIT\_PROCESS 929, 3393
- initstate()* 1063
- INLCR 3211
- inle16()* 1005
- inle32()* 1009
- inodes
  - getting 2417
  - number of 715, 3002
- INPCK 3211
- input, formatted 703, 746, 2533,
  - 3421, 3426, 3430, 3454,
  - 3561
- input\_line()* 1066
- \_input\_line\_max* 1066
- instruction set architecture 318
- instrumented kernel 1093, 3331
- \_\_INT\_BITS\_\_** 111
- integers
  - absolute value 115
  - atomic operations
    - addition 178, 180
    - subtraction 190, 192
  - division 408
  - pseudo-random numbers 2286,
    - 2288
  - quotient 408

- remainder 408
- rounding 228, 639
- size of 111
- system message log, writing
  - to 2873
- integral logarithms 999
- integral part of a floating-point
  - number 1615
- integral power of 2 701, 1428
- Intel 80x86-specific interrupts 1107
- interfaces
  - index, mapping to name 991
  - list of, freeing 989
  - list of, getting 993
  - name, mapping to index 995
- international characters *See* wide characters
- Internet Control Message Protocol
  - See* ICMP
- Internet domain
  - errors 953, 958, 967
  - name servers
    - initializing 2374
    - querying 2378, 2381, 2384, 2387, 2389
  - names
    - compressing 427
    - expanding 429
- Internet Protocol *See* IP
- InterruptAttach()*,
  - InterruptAttach\_r()* 1069
- InterruptAttachEvent()*,
  - InterruptAttachEvent\_r()* 1077
- InterruptDetach()*,
  - InterruptDetach\_r()* 1083
- InterruptDisable()* 1085
- InterruptEnable()* 1087
- InterruptHookIdle()* 1089
- InterruptHookTrace()* 1093
- InterruptLock()* 1095
- InterruptMask()* 1097
- interrupts
  - classes 1069
  - disabling 1085
  - enabling 1087
  - events
    - attaching 1077
    - detaching 1083
  - handlers
    - attaching 1069
    - detaching 1083
    - disabling hardware
      - interrupts 1085
    - guidelines for writing 1072
    - idle, attaching 1089
    - locking 1095, 1100
    - stack size 1072
  - Intel 80x86-specific 1107
  - level-sensitive 1097
  - masking 1097, 1102
  - PCI
    - mapping 1919
    - routing information 1916
  - requests, managing 2494
  - waiting for 1104
- InterruptUnlock()* 1100
- InterruptUnmask()* 1102
- InterruptWait()*, *InterruptWait\_r()* 1104
- \_intr\_v86()* 1107
- inverse hyperbolic cosines 125
- inverse hyperbolic sines 156
- inverse hyperbolic tangents 165

io\_chmod\_t 1146  
 io\_chown\_t 1151  
 io\_close\_t 1158  
 io\_devctl\_t 1167  
 io\_fdinfo\_t 1176  
 io\_link\_t 1183  
 io\_lseek\_t 1198  
 io\_mknod\_t 1203  
 io\_mmap\_t 1206  
 io\_notify\_t 1213  
 io\_open\_t 1238  
 io\_openfd\_t 1244  
 io\_pathconf\_t 1250  
 io\_read\_t 1257  
 io\_readlink\_t 1261  
 io\_rename\_t 1264  
 io\_space\_t 1268  
 io\_stat\_t 1274  
 io\_sync\_t 1280  
 io\_unlink\_t 1291  
 io\_utime\_t 1296  
 io\_write\_t 1304  
 \_IO\_CHMOD 1145, 1148  
 \_IO\_CHOWN 1150, 1153  
 \_IO\_CLOSE 1157, 1160  
 \_IO\_COMBINE\_FLAG 1146, 1151,  
     1158, 1168, 1176, 1198,  
     1206, 1214, 1244, 1250,  
     1258, 1268, 1274, 1281,  
     1296, 1305  
 IO\_CONNECT 1111  
 \_IO\_CONNECT 1241  
 \_IO\_CONNECT\_COMBINE 1111  
 \_IO\_CONNECT\_COMBINE\_CLOSE 1111  
 \_IO\_CONNECT\_EFLAG\_DIR 1116,  
     1121  
 \_IO\_CONNECT\_EFLAG\_DOT 1116,  
     1121  
 \_IO\_CONNECT\_EXTRA\_LINK 1116  
 \_IO\_CONNECT\_EXTRA\_MOUNT 1116  
 \_IO\_CONNECT\_EXTRA\_MOUNT\_OCB  
     1116  
 \_IO\_CONNECT\_EXTRA\_MQUEUE 1116  
 \_IO\_CONNECT\_EXTRA\_NONE 1116  
 \_IO\_CONNECT\_EXTRA\_PHOTON 1116  
 \_IO\_CONNECT\_EXTRA\_PROC\_SYMLINK  
     1116  
 \_IO\_CONNECT\_EXTRA\_RENAME 1116  
 \_IO\_CONNECT\_EXTRA\_RESMGR\_LINK  
     1116  
 \_IO\_CONNECT\_EXTRA\_SEM 1116  
 \_IO\_CONNECT\_EXTRA\_SOCKET 1116  
 \_IO\_CONNECT\_EXTRA\_SYMLINK 1116  
 \_IO\_CONNECT\_LINK 1112  
 \_IO\_CONNECT\_MKNOD 1112  
 \_IO\_CONNECT\_MOUNT 1112  
 \_IO\_CONNECT\_OPEN 1111  
 \_IO\_CONNECT\_READLINK 1112  
 \_IO\_CONNECT\_RENAME 1112,  
     1264  
 \_IO\_CONNECT\_RET\_LINK 1237  
 \_IO\_CONNECT\_RSVD\_UNBLOCK 1112  
 \_IO\_CONNECT\_UNLINK 1111  
 ioctl() 1123  
 \_IO\_DEVCTL 1166, 1170  
 \_IOFBF 2742  
 \_IO\_FDINFO 1173, 1175  
 iofdinfo() 1125  
 \_IO\_FLAG\_RD 1276  
 \_IO\_FLAG\_WR 1276  
 iofunc\_attr\_t 1132  
 iofunc\_notify\_t 1215  
 iofunc\_ocb\_t 1234

- IOFUNC\_ATTR\_ETIME 1133, 1283, 1299
- IOFUNC\_ATTR\_CTIME 1133, 1283, 1299
- IOFUNC\_ATTR\_DIRTY\_MODE 1133, 1299
- IOFUNC\_ATTR\_DIRTY\_MTIME 1133
- IOFUNC\_ATTR\_DIRTY\_NLINK 1133
- IOFUNC\_ATTR\_DIRTY\_OWNER 1133
- IOFUNC\_ATTR\_DIRTY\_RDEV 1133
- IOFUNC\_ATTR\_DIRTY\_SIZE 1133
- IOFUNC\_ATTR\_DIRTY\_TIME 1133, 1283, 1299
- iofunc\_attr\_init()* 1127
- iofunc\_attr\_lock()* 1129
- IOFUNC\_ATTR\_MTIME 1283, 1299
- IOFUNC\_ATTR\_PRIVATE 1133
- IOFUNC\_ATTR\_T 1234
- iofunc\_attr\_trylock()* 1137
- iofunc\_attr\_unlock()* 1139
- iofunc\_check\_access()* 1141
- iofunc\_chmod()* 1145
- iofunc\_chmod\_default()* 1148
- iofunc\_chown()* 1150
- iofunc\_chown\_default()* 1153
- iofunc\_client\_info()* 1155
- iofunc\_close\_dup()* 1157
- iofunc\_close\_dup\_default()* 1160
- iofunc\_close\_ocb()* 1162
- iofunc\_close\_ocb\_default()* 1164
- iofunc\_devctl()* 1166
- iofunc\_devctl\_default()* 1170
- iofunc\_fdinfo()* 1173
- iofunc\_fdinfo\_default()* 1175
- iofunc\_func\_init()* 1179
- iofunc\_link()* 1182
- iofunc\_lock()* 1186
- iofunc\_lock\_calloc()* 1188
- iofunc\_lock\_default()* 1190
- iofunc\_lock\_free()* 1193
- iofunc\_lock\_ocb\_default()* 1195
- iofunc\_lseek()* 1197
- iofunc\_lseek\_default()* 1200
- iofunc\_mknod()* 1202
- iofunc\_mmap()* 1205
- iofunc\_mmap\_default()* 1209
- IOFUNC\_MOUNT\_32BIT 1197
- IOFUNC\_MOUNT\_T 1132
- iofunc\_notify()* 1213
- IOFUNC\_NOTIFY\_INPUT 1220
- IOFUNC\_NOTIFY\_OBAND 1220
- IOFUNC\_NOTIFY\_OUTPUT 1220
- iofunc\_notify\_remove()* 1218
- iofunc\_notify\_trigger()* 1220
- iofunc\_ocb\_attach()* 1224
- iofunc\_ocb\_calloc()* 1225
- iofunc\_ocb\_detach()* 1228
- IOFUNC\_OCB\_FLAGS\_PRIVATE 1235
- iofunc\_ocb\_free()* 1231
- IOFUNC\_OCB\_PRIVILEGED 1235
- iofunc\_open()* 1236
- iofunc\_open\_default()* 1241
- iofunc\_openfd()* 1243
- iofunc\_openfd\_default()* 1247
- iofunc\_pathconf()* 1249
- iofunc\_pathconf\_default()* 1252
- iofunc\_read\_default()* 1254
- iofunc\_readlink()* 1260
- iofunc\_read\_verify()* 1256
- iofunc\_rename()* 1264
- iofunc\_space\_verify()* 1267
- iofunc\_stat()* 1271
- iofunc\_stat\_default()* 1273
- iofunc\_sync()* 1276

- iofunc\_sync\_default()* 1278
- iofunc\_sync\_verify()* 1280
- iofunc\_time\_update()* 1283
- iofunc\_unblock()* 1285
- iofunc\_unblock\_default()* 1287
- iofunc\_unlink()* 1290
- iofunc\_unlock\_ocb\_default()* 1293
- iofunc\_utime()* 1295
- iofunc\_utime\_default()* 1298
- iofunc\_write\_default()* 1301
- iofunc\_write\_verify()* 1303
- \_IOLBF 2742
- \_IO\_LOCK 1190
- \_IO\_LSEEK 1197, 1200
- \_IO\_MMAP 1205, 1209
- \_IONBF 2742
- ionotify()* 1307
- \_IO\_OPENFD 1243, 1247
- \_IO\_OPENFD\_NONE 1245
- \_IO\_OPENFD\_PIPE 1245
- \_IO\_OPENFD\_RESERVED 1245
- \_IO\_PATHCONF 1249, 1252
- \_IO\_READ 1254, 1256
- IO\_SET\_CONNECT\_RET()* 1237
- IO\_SET\_PATHCONF\_VALUE()* 1249
- IO\_SET\_READ\_NBYTES()* 1257
- IO\_SET\_WRITE\_NBYTES()* 1304
- \_IO\_SPACE 1267
- \_IO\_STAT 1271, 1273
- \_IO\_SYNC 1278, 1280
- \_IO\_UTIME 1296, 1298
- iov\_t* 2334, 2686, 3558
  - base, getting 821
  - fields, filling 2686
  - length, getting 823
  - reading from a file 2334
  - resource managers 2446
  - writing to a file 3558
- iovec**
  - copying 1540
- \_IO\_WRITE 1301
- \_IO\_XFLAG\_BLOCK 1258, 1305
- \_IO\_XFLAG\_DIR\_EXTRA\_HINT 1258, 1305
- \_IO\_XFLAG\_NONBLOCK 1258, 1305
- \_IO\_XTYPE\_MQUEUE 1258, 1305
- \_IO\_XTYPE\_NONE 1258, 1305
- \_IO\_XTYPE\_OFFSET 1258, 1305
- \_IO\_XTYPE\_READCOND 1258, 1305
- \_IO\_XTYPE\_REGISTRY 1258, 1305
- \_IO\_XTYPE\_TCPIP 1258, 1305
- \_IO\_XTYPE\_TCPIP\_MSG 1258, 1305
- IP (Internet Protocol) 1313
- ip\_mreq** 1317
- IP addresses
  - CIDR (Classless Internet Domain Routing),
    - converting to/from 1023, 1028
  - classes of 1024
  - extracting
    - network number 1026
  - local network addresses,
    - extracting 1019
  - manipulating 1021
  - network numbers, converting
    - to/from strings 1030
  - specifying in dot notation 1024
  - strings, converting
    - to/from 1015, 1017, 1032, 1034, 1036, 1039

- IP Payload Compression
  - Protocol 1325
- IP6 (Internet Protocol v6) 1338
- IP\_ADD\_MEMBERSHIP 1317
- IP\_DROP\_MEMBERSHIP 1317
- IP\_HDRINCL 903, 1314
- IP\_MAX\_MEMBERSHIPS 1317
- IP\_MULTICAST\_IF 1316
- IP\_MULTICAST\_LOOP 1316
- IP\_MULTICAST\_TTL 1315
- IP\_OPTIONS 1313
- IPPROTO\_ICMPV6 985, 986
- IPPROTO\_RAW 1314, 1345
- IP\_RECVDSTADDR 1314
- IP\_RECVIF 1314
- IPsec (secure Internet Protocol) 1320
- IP\_TOS 903, 1313
- IP\_TTL 1313
- ipv6\_mreq** 1340
- IPV6\_BINDV6ONLY 1341
- IPV6\_DSTOPTS 1342
- IPV6\_HOPLIMIT 1342
- IPV6\_HOPOPTS 1342
- IPV6\_JOIN\_GROUP 1340
- IPV6\_MULTICAST\_HOPS 1339
- IPV6\_MULTICAST\_IF 1339
- IPV6\_MULTICAST\_LOOP 1339
- IPV6\_PKTINFO 1341
- IPV6\_PORTRANGE 1341
- IPV6\_RTHDR 1342
- IPV6\_UNICAST\_HOPS 1338
- IRQs (Interrupt ReQuests),
  - managing 2494
- isalnum()* 1348
- isalpha()* 1350
- isascii()* 1352
- isatty()* 1354
- iscntrl()* 1356
- isdigit()* 1358
- isfdtype()* 1360
- isgraph()* 1362
- ISIG 3213
- isinf()*, *isinff()* 1364
- islower()* 1366
- isnan()*, *isnanf()* 1368
- isprint()* 1370
- ispunct()* 1372
- isspace()* 1374
- ISTRIP 3211
- isupper()* 1377
- iswalnum()* 1379
- iswalpha()* 1381
- iswcntrl()* 1383
- iswctype()* 1385
- iswdigit()* 1387
- iswgraph()* 1389
- iswlower()* 1391
- iswprint()* 1393
- iswpunct()* 1395
- iswspace()* 1397
- iswupper()* 1399
- iswxdigit()* 1401
- isxdigit()* 1403
- ITIMER\_REAL 826, 2689
- itimerspec** 3269
- itoa()* 1405
- IXOFF 3211
- IXON 3211



**J**

*j0()*, *j0f()* 1408  
*j1()*, *j1f()* 1410  
*jn()*, *jnf()* 1412  
 job control, supporting 3137  
*jrands48()* 1414  
 jumps, nonlocal 1473, 2691, 2787,  
 2832

**K**

**kerinfo** 3152  
 kernel  
   blocking states, setting timeouts  
     on 3299  
   calls  
     *ChannelCreate()*,  
       *ChannelCreate\_r()* 248  
     *ChannelDestroy()*,  
       *ChannelDestroy\_r()* 255  
     *ClockAdjust()*,  
       *ClockAdjust\_r()* 282  
     *ClockCycles()* 284  
     *ClockId()*, *ClockId\_r()* 300  
     *ClockPeriod()*,  
       *ClockPeriod\_r()* 304  
     *ClockTime()*, *ClockTime\_r()*  
       306  
     *ConnectAttach()*,  
       *ConnectAttach\_r()* 327  
     *ConnectClientInfo()*,  
       *ConnectClientInfo\_r()*  
       331

*ConnectDetach()*,  
       *ConnectDetach\_r()* 251,  
       335  
     *ConnectFlags()*,  
       *ConnectFlags\_r()* 337  
     *ConnectServerInfo()*,  
       *ConnectServerInfo\_r()*  
       340  
     *DebugBreak()* 363  
     *DebugKDBreak()* 365  
     *DebugKDOutput()* 366  
     *InterruptAttach()*,  
       *InterruptAttach\_r()* 1069  
     *InterruptAttachEvent()*,  
       *InterruptAttachEvent\_r()*  
       1077  
     *InterruptDetach()*,  
       *InterruptDetach\_r()* 1083  
     *InterruptDisable()* 1085  
     *InterruptEnable()* 1087  
     *InterruptHookIdle()* 1089  
     *InterruptHookTrace()* 1093  
     *InterruptLock()* 1095  
     *InterruptMask()* 1097  
     *InterruptUnlock()* 1100  
     *InterruptUnmask()* 1102  
     *InterruptWait()*,  
       *InterruptWait\_r()* 1104  
     *MsgDeliverEvent()*,  
       *MsgDeliverEvent\_r()*  
       1661  
     *MsgError()*, *MsgError\_r()*  
       1669  
     *MsgInfo()*, *MsgInfo\_r()*  
       1672  
     *MsgKeyData()*,  
       *MsgKeyData\_r()* 1675

- MsgRead()*, *MsgRead\_r()*  
1682
- MsgReady()*, *MsgReady\_r()*  
1686
- MsgReceive()*,  
*MsgReceive\_r()* 1689
- MsgReceivePulse()*,  
*MsgReceivePulse\_r()*  
1694
- MsgReceivePulsev()*,  
*MsgReceivePulsev\_r()*  
1697
- MsgReceivev()*,  
*MsgReceivev\_r()* 1700
- MsgReply()*, *MsgReply\_r()*  
1704
- MsgReplyv()*, *MsgReplyv\_r()*  
1707
- MsgSend()*, *MsgSend\_r()*  
1710
- MsgSendnc()*, *MsgSendnc\_r()*  
1714
- MsgSendPulse()*,  
*MsgSendPulse\_r()* 1719
- MsgSendsv()*, *MsgSendsv\_r()*  
1722
- MsgSendsvnc()*,  
*MsgSendsvnc\_r()* 1726
- MsgSendv()*, *MsgSendv\_r()*  
1730
- MsgSendvnc()*,  
*MsgSendvnc\_r()* 1734
- MsgSendvs()*, *MsgSendvs\_r()*  
1738
- MsgSendvsnc()*,  
*MsgSendvsnc\_r()* 1742
- MsgVerifyEvent()*,  
*MsgVerifyEvent\_r()* 1746
- MsgWrite()*, *MsgWrite\_r()*  
1748
- MsgWritev()*, *MsgWritev\_r()*  
1752
- SchedGet()*, *SchedGet\_r()*  
2570
- SchedInfo()*, *SchedInfo\_r()*  
2574
- SchedSet()*, *SchedSet\_r()*  
2577
- SchedYield()*, *SchedYield\_r()*  
2579
- SignalAction()*,  
*SignalAction\_r()* 2796
- SignalKill()*, *SignalKill\_r()*  
2804
- SignalProcmask()*,  
*SignalProcmask\_r()* 2810
- SignalSuspend()*,  
*SignalSuspend\_r()* 2813
- SignalWaitinfo()*,  
*SignalWaitinfo\_r()* 2816
- SyncCondvarSignal()*,  
*SyncCondvarSignal\_r()*  
3104
- SyncCondvarWait()*,  
*SyncCondvarWait\_r()*  
3107
- SyncCtl()*, *SyncCtl\_r()* 3112
- SyncDestroy()*,  
*SyncDestroy\_r()* 3114
- SyncMutexEvent()*,  
*SyncMutexEvent\_r()* 3117
- SyncMutexLock()*,  
*SyncMutexLock\_r()* 3119

*SyncMutexRevive()*,  
     *SyncMutexRevive\_r()*  
     3122  
*SyncMutexUnlock()*,  
     *SyncMutexUnlock\_r()*  
     3124  
*SyncSemPost()*,  
     *SyncSemPost\_r()* 3127  
*SyncSemWait()*,  
     *SyncSemWait\_r()* 3129  
*SyncTypeCreate()*,  
     *SyncTypeCreate\_r()* 3133  
*ThreadCancel()*,  
     *ThreadCancel\_r()* 3234  
*ThreadCreate()*,  
     *ThreadCreate\_r()* 3239  
*ThreadCtl()*, *ThreadCtl\_r()*  
     1069, 1078, 1083, 1085,  
     1087, 1095, 1097, 1100,  
     1102, 3245  
*ThreadDestroy()*,  
     *ThreadDestroy\_r()* 3249  
*ThreadDetach()*,  
     *ThreadDetach\_r()* 3252  
*ThreadJoin()*, *ThreadJoin\_r()*  
     3254  
*TimerAlarm()* 3282  
*TimerCreate()* 3285  
*TimerDestroy()*,  
     *TimerDestroy\_r()* 3288  
*TimerInfo()*, *TimerInfo\_r()*  
     3291  
*TimerSettime()*,  
     *TimerSettime\_r()* 3295  
*TimerTimeout()*,  
     *TimerTimeout\_r()* 3299  
*TraceEvent()* 3331

debugging 365, 366  
 instrumented 1093, 3331  
 thread scheduler 2574  
*kill()* 1416  
*killpg()* 1419

## L

*labs()* 1421  
 large-file support 103  
     file information 710, 1485,  
     2993  
     filesystem information 714,  
     3001  
     mapped memory, offset  
     of 1527  
     memory, offset of 1959  
     opening 349, 1835  
     position  
     setting 1482, 3204  
     reading 1965  
     shared memory, mapping 1585  
     symbolic link  
     information 1485  
     system-resource limits 881,  
     2719  
     truncating 726  
     writing 2272  
 LC\_ALL 2698  
 LC\_COLLATE 2698  
 LC\_CTYPE 1385, 2698  
*lchown()* 1423  
 LC\_MESSAGES 2698  
 LC\_MONETARY 2698  
 LC\_NUMERIC 2698

- lcong48()* 1426
- lconv** 1449
- L\_ctermid 355
- LC\_TIME 2698
- ldexp()*, *ldexpf()* 1428
- ldiv\_t** 1430
- ldiv()* 1430
- LD\_LIBRARY\_PATH** 422
- length, calculating
  - hypotenuse 981
  - strings 3039
  - wide-character strings 3490
- lesser of two numbers 1561
- level-sensitive interrupts 1097
- lfind()* 1433
- lgamma()*, *lgamma\_r()*, *lgammaf()*,  
*lgammaf\_r()* 1435
- LIB** 2602
- libraries, locating 319
- limits
  - core files, size of 2720
  - data segment, size of 2720
  - device numbers 2510
  - files
    - descriptors, number of 2720
    - link count 673, 1866
    - maximum per process 3137
    - names, length of 673, 1866
    - size 2720
  - filesystems 714, 3001
  - hops 1338, 1342
  - host names, length 2684
  - inheriting 519
  - iov* arrays 2334
  - path names, length of 674,  
1867
- pipes, number of bytes written
  - atomically 674, 1867
- processes
  - argument lists 3136
  - CPU time 2720
  - execution time 2559
  - files, number open 3137
  - I/O requests 1442
  - mapped address space 2720
  - maximum per real user
    - ID 3136
  - scheduling policy 2547,  
2549
  - supplementary group
    - IDs 3136
- sockets, pending
  - connections 1447, 2475
- stack size 2720
- system resources
  - getting 773, 881
  - setting 2719
- TCP maximum segment
  - size 3190
- terminals
  - canonical input buffer
    - size 673, 1866
  - raw input buffer size 673,  
1866
- threads
  - execution time 2574
  - priority 2574
  - stack size 2039
- \_\_LINE\_\_** 158
- line buffering, setting for stream
  - I/O 2696
- linear search 1433, 1479
- LINES** 274

- link()* 1438
  - resource managers,
    - implementing in 1182
- link-local addresses 1052
- linker symbols
  - \_btext* 217
  - \_edata* 459
  - \_end* 462
  - \_etext* 515
- LINK\_MAX 632, 1439
- links, symbolic
  - creating 3099
  - deleting 2368, 3380
  - information, getting 1485
  - ownership, changing 1423
  - reading 2331
  - resolving 2341
  - temporary 1874, 1876
- lio\_listio()* 1443
- LIO\_NOP 1443
- LIO\_NOWAIT 1442
- LIO\_READ 1443
- LIO\_WAIT 1442
- LIO\_WRITE 1443
- listen()* 1447, 3189, 3573
- little endian
  - big endian, converting
    - to/from 477, 479, 481, 483, 485, 487, 3093
  - `__LITTLEENDIAN__` manifest 111
  - messages 1550
  - native format, converting
    - to/from 471, 473, 475
  - ports
    - reading from 1005, 1009
    - writing to 1858, 1862
  - unaligned values
    - accessing safely 3364, 3366, 3368
    - writing safely 3358, 3360, 3362
- `__LITTLEENDIAN__` 111
- lltoa()* 1488
- local network addresses, converting
  - to/from IP addresses 1021
- local times, converting to/from
  - calendar times 1454, 1456, 1577
- LOCAL\_CREDS 3379
- LOCALDOMAIN** 2374, 2378, 2382, 2384, 2387, 2390
- localeconv()* 1449
- locales
  - classes, wide-character 3533
  - daylight saving time 362
  - numeric formatting 1449
  - setting 2699
  - strings, comparing 3021, 3482
- localtime()* 1454
- localtime\_r()* 1456
- LOCK\_EX 634
- lockf()* 1459
- LOCK\_NB 634
- locks
  - files 589, 634, 637, 729, 1459
  - mutexes
    - attributes, destroying 2137
    - attributes, initializing 2150
    - attributes, priority ceiling 2139, 2152

- attributes,
    - process-shared 2143, 2156
  - attributes, recursive 2145, 2158
  - attributes, scheduling
    - protocol 2141, 2154
  - attributes, type 2147, 2161
  - destroying 2118, 3114
  - events 3112, 3117
  - initializing 2122, 3133
  - locking 2124, 2130, 2133, 3119, 3124
  - priority 3112
  - priority ceiling 2120, 2128
  - reviving 3122
  - unlocking 2135
- read-write
- attributes, creating 2191
  - attributes, destroying 2187
  - attributes,
    - process-shared 2189, 2193
  - destroying 2166, 3114
  - initializing 2168, 3133
  - locking for reading 2171, 2173, 2179
  - locking for writing 2176, 2181, 2185
  - unlocking 2183
- sleepon
- destroying 2855
  - initializing 2857
  - locking 2210, 2859
  - unblocking 2208, 2212, 2853, 2861
  - unlocking 2218, 2863
  - waiting 2214, 2220, 2865
- LOCK\_SH 634
- LOCK\_UN 634
- log()*, *logf()* 1462
- log, system message
  - closing 314
  - log priority mask 2701
  - opening 1850
  - writing to 3147, 3448
    - blocks 2867
    - formatted output 2869, 3433
    - integers 2873
- log10()*, *log10f()* 1466
- log1p()*, *log1pf()* 1464
- LOG\_ALERT 3148
- logarithms
  - base 10 1466
  - integral 999
  - natural 1462
  - $x + 1$  1464
- logb()*, *logbf()* 1468
- LOG\_CRIT 3148
- LOG\_DEBUG 3148
- LOG\_EMERG 3148
- LOG\_ERR 3148
- logging in
  - previous lines, discarding 1602
  - pseudo-ttys 659, 1471
- logical interrupt vector
  - numbers 1069
- LOG\_INFO 3148
- LOGIN\_PROCESS 929, 932, 3393
- login\_tty()* 1471
- LOG\_MASK()* 2701
- LOG\_NOTICE 3148
- LOG\_UPTO()* 2701
- LOG\_WARNING 3148

- long integers
    - absolute value 1421
    - division 1430
    - pseudo-random numbers
      - nonnegative 1476, 1825
      - signed 1414, 1656, 2290
    - quotient 1430
    - remainder 1430
    - size of 111
  - `__LONG_BITS__` 111
  - `longjmp()` 1473, 2283
  - lowercase
    - characters, converting to 3321
    - strings, converting to 3041
    - testing a character for 1366, 1391
    - wide characters, converting to 3325, 3327, 3531
  - `lrand48()` 1476
  - `lsearch()` 1479
  - `lseek()`, `lseek64()` 1482
    - resource managers, implementing in 1197, 1200
  - `lstat()`, `lstat64()` 1485
  - `L_tmpnam` 3318
  - `ltoa()` 1488
  - `ltrunc()` 1491
- M**
- Machine Status Register 3152
  - `main()` 1495, 2791
    - arguments
      - auxiliary 198
      - number of 150
      - parsing 842
      - vector of 151
  - major device numbers 2417, 2510, 2997
  - `major()` 2997
  - `_MAJOR_BLK_PREFIX` 2509
  - `_MAJOR_CHAR_PREFIX` 2509
  - `_MAJOR_DEV` 2509
  - `_MAJOR_FSYS` 2509
  - `_MAJOR_PATHMGR` 2509
  - `makedev()` 2997
  - `mallinfo` 1498
  - `mallinfo()` 1498
  - `malloc()` 1500
  - `MALLOC_ARENA_SIZE` 1502
  - `MALLOC_MONOTONIC_GROWTH` 1502
  - `mallopt()` 1503
  - manifests 111
  - `MAP_ANON` 1586
  - `MAP_BELOW16M` 1586
  - `MAP_FAILED` 1588, 1595
  - `MAP_FIXED` 1585, 1586, 1594, 1595
  - `MAP_LAZY` 1586
  - `MAP_NOX64K` 1587
  - `MAP_PHYS` 1587
  - `MAP_PRIVATE` 1585
  - `MAP_SHARED` 1585, 1586
  - `MAP_STACK` 1587
  - `MAP_TYPE` 1585
  - mathematics
    - absolute values 115, 222, 574, 1421
    - Bessel functions 1408, 1410, 1412, 3563, 3565, 3567

- complementary error
  - function 502
- division 408, 1430
- error function 500
- exponentials 569, 571, 701, 1428, 2523, 2526, 2819
- finite numbers 629
- floating-point settings 661, 664, 667, 670
- gamma functions 750, 1435
- hyperbolic functions 125, 156, 165, 347, 2849, 3163
- hypotenuse, length of 981
- infinite numbers 1364
- logarithms 999, 1462, 1464, 1466
- maximum 1504
- minimum 1561
- modular arithmetic 643
- next representable
  - number 1816
- not a number, determining
  - if 1368
- powers 1963
- pseudo-random numbers 431, 498, 1063, 1414, 1426, 1476, 1656, 1825, 2286, 2288, 2290, 2605, 2733, 2982, 2984, 2986
- radix-independent
  - exponents 1468
- remainders 433, 2366
- roots 226, 2980
- rounding 228, 639, 2472
- sign, copying 343
- trigonometry 123, 154, 161, 163, 345, 2847, 3161
- max()* 1504
- MAXHOSTNAMELEN 2684
- Maximum Segment Size (MSS) 3190
- MB\_CUR\_MAX 3528
- mblen()* 1506
- mbrlen()* 1509
- mbrtowc()* 1511
- mbsinit()* 1514
- mbsrtowcs()* 1516
- mbstate\_t** 1514
- mbstowcs()* 1518
- mbtowc()* 1521
- mcheck\_status** 1623
- mcheck()* 1524
- MCHECK\_DISABLED 1623
- MCHECK\_FREE 1624
- MCHECK\_HEAD 1624
- MCHECK\_OK 1624
- MCHECK\_TAIL 1624
- memalign()* 1530
- members, offset of within a
  - structure 1833
- memccpy()* 1532
- memchr()* 1534
- memcmp()* 1536
- memcpy()* 1538
- memcpyv()* 1540
- memicmp()* 1542
- memmove()* 1544
- mem\_offset(), mem\_offset64()* 1527
- memory
  - allocating
    - aligned 1530, 1961
    - \_amblksiz* 2875
  - array 224, 2529
  - automatic (from stack) 144



- blocks 1500, 2338, 2875
  - break value, changing 149, 2520
  - consistency check 1524, 1623
  - controlling 1503
  - data segment, changing 211
  - heap block, aligned on page boundary 3412
  - information about, getting 1498
  - comparing 1536, 1542, 3537
  - copying 1532, 1538, 1540, 3539
    - overlapping objects 1544, 3541
  - devices
    - I/O, mapping 1591, 1764
    - physical, mapping into process's address space 1594
  - direct memory access (DMA)
    - channels, managing 2494
  - free, amount of 2999
  - freeing 240, 691, 2338, 2745, 2988
  - locking 1580, 1582, 1758
  - managing 2494
  - mapped
    - contiguous length 1527
    - maximum size 2720
    - offset of 1527
  - offset of, getting 1959
  - physical storage, synchronizing with 1756
  - reallocating 2338, 2988
  - searching
    - for a character 1534
    - for a wide character 3535
  - setting 1546, 3543
  - shared
    - mapping 1585
    - unmapping 1762
  - unlocking 1760
  - unmapping 1766
- memset()* 1546
- message\_attr\_t** 1549
- message queues
  - attributes 1630, 1646
  - closing 1628
  - messages
    - receiving 1640, 1648
    - sending 1643, 1651
  - notifying when nonempty 1633
  - opening 1637
  - persistence of 1638
  - receive-only 1636
  - send-only 1636
  - send-receive 1636
  - unlinking 1654
- message\_attach()* 1549
- message\_connect()* 1555
- message\_detach()* 1558
- messages
  - channels
    - attaching to a process 327
    - creating 248, 1768
    - destroying 255, 1777
  - dispatch interface
    - handlers 1549, 1558
  - errors, handling 1669
  - information about
    - getting 1672

- structure 1658
- Internet domain name servers
  - errors 953, 958, 967
  - queries 2378, 2381, 2384, 2387, 2389
  - sending and
    - interpreting 427, 429, 2374
  - key, adding 1675
  - reading data 1682, 1686, 2438, 2440
  - receiving 1689, 1700
  - replying 1704, 1707
  - resource managers
    - blocking while waiting for 2401
    - handling 2425
    - sending 1710, 1714, 1722, 1726, 1730, 1734, 1738, 1742
  - SNMP
    - creating 2887
    - freeing 2879
    - reading 2889
    - sending 2894
  - sockets
    - peeking at 2343, 2346, 2350
    - receiving from 2344, 2347, 2351
    - sending to 2651, 2653, 2657
    - tampering, preventing 1675
    - unblocking 1669
    - writing data 1748, 1752, 2442, 2444
  - \_MFLAG\_OCB 1617, 1620
  - M\_GRAIN 1502
  - min()* 1561
  - minor device numbers 2417, 2510, 2997
  - minor()* 2997
  - misaligned access response 3245
  - mkdir()* 1563, 3353
  - M\_KEEP 1502
  - mkfifo()* 1566, 3353
  - mknod()* 1133, 1569
    - resource managers, implementing in 1202
  - mkstemp()* 1573
  - mktemp()* 1575
  - mktime()* 1577
  - mlock()* 1580
  - mlockall()* 1582
  - mmap()*, *mmap64()* 1585
    - resource managers, implementing in 1205, 1209
  - mmap\_device\_io()* 1591
  - mmap\_device\_memory()* 1594
  - M\_MMAP\_MAX 1502
  - M\_MMAP\_THRESHOLD 1502
  - M\_MXFAST 1502
  - M\_NLBLKS 1502
  - modem\_script** 1605
  - MODEM\_ALLOW8BIT 1601
  - MODEM\_ALLOWCASE 1601
  - MODEM\_ALLOWCTRL 1601
  - MODEM\_BAUD 1605, 1610
  - MODEM\_LASTLINE 1602
  - MODEM\_NOECHO 1605, 1606, 1610
  - modem\_open()* 1597
  - modem\_read()* 1602
  - modems
    - opening 1597

- reading 1602
- script, running on 1605
  - states 1607
- writing 1612
  - escape characters 1612
  - special characters 1612
- modem\_script()* 1605
- modem\_write()* 1612
- modf()*, *modff()* 1615
- modular arithmetic, floating
  - point 643
- mount()* 1618
- `_MOUNT_AFTER` 1618, 1620
- `_MOUNT_ATIME` 1620
- `_MOUNT_BEFORE` 1617, 1620
- `_MOUNT_CREAT` 1620
- `_MOUNT_ENUMERATE` 1618, 1620
- `_MOUNT_FORCE` 1618, 1620, 3356
- `_MOUNT_NOATIME` 1617, 1620
- `_MOUNT_NOCREATE` 1617, 1620
- `_MOUNT_NOEXEC` 1167, 1617, 1620
- `_MOUNT_NOSUID` 1167, 1617, 1620
- `_MOUNT_OFF32` 1617, 1620
- `_MOUNT_OPAQUE` 1618, 1620
- mount\_parse\_generic\_args()* 1620
- `_MOUNT_READONLY` 1167, 1617, 1620
- `_MOUNT_REMOUNT` 1618, 1620
- `_MOUNT_SUID` 1620
- `_MOUNT_UNMOUNT` 1618, 1620
- mprobe()* 1623
- mprotect()* 1625
- `mq_attr` 1630
- mq\_close()* 1628
- mq\_getattr()* 1630
- mq\_notify()* 1633
- mq\_open()* 1637
- `MQ_PRIO_MAX` 1643
- mq\_receive()* 1640
- mq\_send()* 1643
- mq\_setattr()* 1646
- mq\_timedreceive()* 1648
- mq\_timedsend()* 1651
- mq\_unlink()* 1654
- mrnd48()* 1656
- `MS_ASYNC` 1755
- `MSG_CTRUNC` 2352
- MsgDeliverEvent()*,  
*MsgDeliverEvent\_r()*  
1309, 1661
- `MSG_DONTROUTE` 2650, 2653, 2656
- `MSG_EOR` 2352
- MsgError()*, *MsgError\_r()* 1669
- `MSG_FLAG_CROSS_ENDIAN` 402
- `MSG_FLAG_SIDE_CHANNEL` 1555
- `msghdr` 2351
- MsgInfo()*, *MsgInfo\_r()* 1672
- MsgKeyData()*, *MsgKeyData\_r()*  
1675
- `MSG_OOB` 905, 2343, 2346, 2350, 2352, 2650, 2653, 2656
- `MSG_PEEK` 2343, 2346, 2350
- MsgRead()*, *MsgRead\_r()* 1682
- MsgReadv()*, *MsgReadv\_r()* 1686
- MsgReceive()*, *MsgReceive\_r()*  
1689
- MsgReceivePulse()*,  
*MsgReceivePulse\_r()*  
1694

- MsgReceivePulsev()*,  
    *MsgReceivePulsev\_r()* 1697
  - MsgReceivev()*, *MsgReceivev\_r()* 1700
  - MsgReply()*, *MsgReply\_r()* 1704
  - MsgReplyv()*, *MsgReplyv\_r()* 1707
  - MsgSend()*, *MsgSend\_r()* 1710
  - MsgSendnc()*, *MsgSendnc\_r()* 1714
  - MsgSendPulse()*, *MsgSendPulse\_r()* 1719
  - MsgSendsv()*, *MsgSendsv\_r()* 1722
  - MsgSendsvnc()*, *MsgSendsvnc\_r()* 1726
  - MsgSendv()*, *MsgSendv\_r()* 1730
  - MsgSendvnc()*, *MsgSendvnc\_r()* 1734
  - MsgSendvs()*, *MsgSendvs\_r()* 1738
  - MsgSendvsnc()*, *MsgSendvsnc\_r()* 1742
  - MSG\_TRUNC 2352
  - MsgVerifyEvent()*,  
    *MsgVerifyEvent\_r()* 1746
  - MSG\_WAITALL 2343, 2347, 2350
  - MsgWrite()*, *MsgWrite\_r()* 1748
  - MsgWritev()*, *MsgWritev\_r()* 1752
  - MS\_INVALIDATE 1755
  - MS\_INVALIDATE\_ICACHE 1755
  - MSS (Maximum Segment Size) 3190
  - MS\_SYNC 1755
  - msync()* 1756
  - M\_TOP\_PAD 1502
  - M\_TRIM\_THRESHOLD 1502
  - multibyte characters
    - bytes, counting 1506, 1509
    - wide characters
      - conversion object 1514
      - wide characters, converting
        - to/from 1511, 1516, 1518, 1521, 3474, 3502, 3518, 3528
  - munlock()* 1758
  - munlockall()* 1760
  - munmap()* 1762
  - munmap\_device\_io()* 1764
  - munmap\_device\_memory()* 1766
  - mutexes
    - attributes
      - destroying 2137
      - initializing 2150
      - priority ceiling 2139, 2152
      - process-shared 2143, 2156
      - recursive 2145, 2158
      - scheduling protocol 2141, 2154
      - type 2147, 2161
    - destroying 2118, 3114
    - events, attaching 3112, 3117
    - initializing 2122, 3133
    - locking 2124, 2130, 2133, 3119, 3124
    - priority 3112
      - ceiling 2120, 2128
    - reviving 3122
    - unlocking 2135
- ## N
- name\_attach\_t** 1769
  - name servers
    - errors 953, 958, 967

- initializing 2374
- names
  - compressing 427
  - expanding 429
  - queries 2378, 2381, 2384, 2387, 2389
- name\_attach()* 1768
- NAME\_ATTACH\_FLAG\_GLOBAL 1768
- name\_close()* 1775
- name\_detach()* 1777
- NAME\_FLAG\_ATTACH\_GLOBAL 1768, 1779
- NAME\_FLAG\_DETACH\_SAVEDPP 1777
- NAME\_MAX 1780
- name\_open()* 1779
- names
  - binding to sockets 206, 2293
  - domain
    - setting 2664
  - domain, getting 318, 771
  - host
    - getting 817
    - setting 2684
  - peer, getting 850
  - socket, getting 899, 2468
- NAN (not-a-number) 1368
- nanoseconds
  - busy-waiting for 1784, 1786, 1789, 1791, 1793
  - threads, suspending for 1782
  - timespec**, converting to/from 1827, 3310
- nanosleep()* 1782
- nanospin()* 1784
- nanospin\_calibrate()* 1786
- nanospin\_count()* 1789
- nanospin\_ns()* 1791
- nanospin\_ns\_to\_count()* 1793
- nap()* 1796
- napms()* 1797
- natural logarithms 1462
- nbaconnect()* 1798
- nbaconnect\_result()* 1801
- ND2S\_DIR\_HIDE 1807
- ND2S\_DIR\_SHOW 1807
- ND2S\_DOMAIN\_HIDE 1808
- ND2S\_DOMAIN\_SHOW 1808
- ND2S\_LOCAL\_STR 1808
- ND2S\_NAME\_HIDE 1808
- ND2S\_NAME\_SHOW 1808
- ND2S\_QOS\_HIDE 1808
- ND2S\_QOS\_SHOW 1808
- ND2S\_SEP\_FORCE 1809
- NDEBUG 158
- ND\_LOCAL\_NODE 1803
- ND\_NODE\_CMP()* 1803
- NETDB\_INTERNAL 807, 953, 959
- netent** 1805
- netmgr\_ndtostr()* 1806
- netmgr\_remote\_nd()* 1812
- netmgr\_strtond()* 1814
- network
  - database
    - closing 489
    - entries, getting 836, 838, 840
    - opening 2703
    - structure 1805
  - host entries
    - errors 953, 958, 967
    - getting 800, 804, 806, 810, 812, 815
  - network interface addresses
    - freeing 695

- getting 819
- structure 997
- network numbers
  - IP addresses, converting
    - to/from 1021, 1026
  - strings, converting
    - to/from 1030
- network-byte order
  - host-byte order, converting
    - to/from 969, 971, 1829, 1831
- Neutrino classification 106
- NEW\_TIME 929, 3393
- nextafter()*, *nextafterf()* 1816
- nftw()* 1820
- NGROUPS\_MAX 331, 798, 2680
- nice()* 1823
- NO\_DATA 954, 959
- node descriptors
  - comparing 1803
  - current 3370
  - relative to a remote node 1812
  - strings, converting
    - to/from 1806, 1814
- NOFD 1527, 1585, 1586
- NOFLSH 3213
- nonlocal jumps 1473, 2691, 2787, 2832
- NO\_RECOVERY 954, 959
- normalized fractions 701
- not a number, determining if 1368
- \_NOTIFY\_ACTION\_POLL 1309, 1310
- \_NOTIFY\_ACTION\_POLLARM 1310
- \_NOTIFY\_ACTION\_TRANARM 1310
- \_NOTIFY\_COND\_INPUT 1214, 1308
- \_NOTIFY\_COND\_MASK 1308
- \_NOTIFY\_COND\_OBAND 1215, 1308
- \_NOTIFY\_COND\_OUTPUT 1214, 1308
- \_NOTIFY\_DATA\_MASK 1309
- nrand48()* 1825
- nsec2timespec()* 1827
- NSIG 2780
- \_NTO\_CHF\_COID\_DISCONNECT 250
- \_NTO\_CHF\_DISCONNECT 251
- \_NTO\_CHF\_FIXED\_PRIORITY 249, 251
- \_NTO\_CHF\_NET\_MSG 251
- \_NTO\_CHF\_REPLY\_LEN 251, 1659
- \_NTO\_CHF\_SENDER\_LEN 252, 1659
- \_NTO\_CHF\_THREAD\_DEATH 252
- \_NTO\_CHF\_UNBLOCK 252, 253, 1659
- \_NTO\_COF\_CLOEXEC 327, 337
- ntohl()* 1829
- ntohs()* 1831
- \_NTO\_INTR\_CLASS\_EXTERNAL 1069
- \_NTO\_INTR\_CLASS\_SYNTHETIC 1069
- \_NTO\_INTR\_FLAGS\_END 1073, 1074, 1079, 1090
- \_NTO\_INTR\_FLAGS\_PROCESS 1073, 1074, 1079, 1090, 1091
- \_NTO\_INTR\_FLAGS\_TRK\_MSK 1073, 1090
- \_NTO\_INTR\_SPARE 1069
- \_NTO\_KEYDATA\_CALCULATE 1674, 1677
- \_NTO\_KEYDATA\_VERIFY 1674, 1677
- \_NTO\_MI\_UNBLOCK\_REQ 1659

- \_NTO\_RESET\_OVERRUNS 3290
  - \_NTO\_SCTL\_GETPRIOCEILING 3111
  - \_NTO\_SCTL\_SETEVENT 3111
  - \_NTO\_SCTL\_SETPRIOCEILING 3111
  - \_NTO\_SYNC\_COND 3132
  - \_NTO\_SYNC\_MUTEX\_FREE 3132
  - \_NTO\_SYNC\_SEM 3132
  - \_NTO\_TCTL\_ALIGN\_FAULT 3245
  - \_NTO\_TCTL\_IO 1069, 1078, 1083, 1085, 1087, 1095, 1097, 1100, 1102, 3246
  - \_NTO\_TCTL\_RUNMASK 3246
  - \_NTO\_TCTL\_THREADS\_CONT 3247
  - \_NTO\_TCTL\_THREADS\_HOLD 3247
  - \_NTO\_TIMEOUT\_CONDVAR 3300
  - \_NTO\_TIMEOUT\_INTR 3300
  - \_NTO\_TIMEOUT\_JOIN 3300
  - \_NTO\_TIMEOUT\_MUTEX 3300
  - \_NTO\_TIMEOUT\_RECEIVE 3300
  - \_NTO\_TIMEOUT\_REPLY 3300
  - \_NTO\_TIMEOUT\_SEM 3300
  - \_NTO\_TIMEOUT\_SEND 3300
  - \_NTO\_TIMEOUT\_SIGSUSPEND 3300
  - \_NTO\_TIMEOUT\_SIGWAITINFO 3300
  - \_NTO\_TIMER\_SEARCH 3290
  - \_NTO\_TRACE\_\* 3332
  - \_NTO\_VERSION 111
  - numbers
    - determining if
      - finite 629
      - infinite 1364
      - not a number 1368
    - formatting 1449
    - maximum 1504
    - minimum 1561
    - next representable
      - floating-point 1816
    - strings, converting
      - to/from 170, 172, 174, 176, 1405, 1488, 3072, 3076, 3083, 3086, 3350, 3397
      - wide-character strings,
        - converting to/from 3508, 3512, 3516, 3522
- O**
- O\_APPEND 589, 1113, 1245, 1836, 1846, 2917, 2922, 3550
  - OCBs (Open Control Blocks)
    - allocating 1225
    - attaching 1224, 2450
    - detaching 1162, 1164, 1228, 2460
    - freeing 1164, 1231
    - getting 2448
    - structure 1234
    - unlocking 1293
  - O\_CLOEXEC 1113, 1836
  - O\_CREAT 349, 1636, 1836, 2753, 2917, 2918
  - O\_DSYNC 1113, 1276, 1281, 1837, 3556
  - O\_EXCL 1113, 1636, 1837, 1891, 2753, 2917
  - off\_t*, limiting to 32 bits 1617
  - offsetof()* 1833
  - OHFLOW 3212
  - O\_LARGEFILE 1113, 1837
  - OLD\_TIME 929, 3393
  - once-initialization 2163

- O\_NOCTTY 1114, 1837
- O\_NONBLOCK 589, 1256, 1267, 1303, 1603, 1636, 1837, 1947, 2307, 2322, 2335, 3551
- OOB (out-of-band) data
  - determining if at mark 2906
  - sending/receiving 2343, 2346, 2350
- Open Control Blocks *See* OCBs
- open()*, *open64()* 349, 1835, 3353
  - resource managers, implementing in 1241
- opendir()* 1843
- openfd()* 1846
  - resource managers, implementing in 1243, 1247
- openlog()* 1850
- OPEN\_MAX 2723, 2925, 2952
- openpty()* 1852
- operating system
  - name 319, 3370
  - release level 319
  - target 111
  - version 111, 319, 3370
- OPOST 3212
- optarg* 842
- opterr* 843
- optimization, compiling with 111
- \_\_OPTIMIZE\_\_ 111
- optind* 842
- options
  - command-line
    - parsing 842, 918
  - mount, parsing 1620
  - socket-level 902, 2730
- optopt* 843
- O\_RDONLY 589, 1113, 1244, 1636, 1836, 1846, 2753, 2917, 2922
- O\_RDWR 589, 1113, 1244, 1491, 1636, 1836, 1846, 1891, 2753, 2917, 2922
- O\_REALIDS 1114, 1155, 1838
- O\_RSynchronize 1114, 1281, 1838
- O\_SYNC 1114, 1276, 1281, 1839, 3556
- other scheduling 2031
- O\_TRUNC 349, 1114, 1245, 1839, 1846, 2753, 2917, 2922
- out-of-band (OOB) data
  - determining if at mark 2906
  - sending/receiving 2343, 2346, 2350
- out16()* 1858
- out16s()* 1860
- out32()* 1862
- out32s()* 1864
- out8()* 1854
- out8s()* 1856
- outbe16()* 1858
- outbe32()* 1862
- outle16()* 1858
- outle32()* 1862
- output, formatted 676, 741, 1968, 3418, 3424, 3428, 3452, 3548
- overlapping memory, copying 1544, 3541
- ownership, changing of a file 265, 581, 1423



**O\_WRONLY** 349, 589, 1113, 1244,  
1491, 1636, 1836, 1846,  
2917, 2922

## P

packets *See also* ROUTE

routing 2481

SNMP

reading 2889

**P\_ALL** 3466

**PARENB** 3212

**PARMRK** 3211

**PARODD** 3212

**PARSTK** 3212

**passwd** 868, 871, 873, 876, 878

passwords

database

closing 491

entries, getting for a

user 871, 873, 876, 878

entries, getting next 868

rewinding 2713

encrypting 353, 2275

prompting for and reading 848

shadow database

closing 493

entry, reading 618, 911, 915

entry, structure 2260

entry, writing 2260

rewinding 2732

**PATH** 57, 274, 536, 559, 2602,  
2943, 2947, 2952, 2969,  
2973

*pathconf()* 1866

resource managers,

implementing in 1249,

1252

*pathfind()*, *pathfind\_r()* 1871

**PATH\_MAX** 768, 1780

*pathmgr\_symlink()* 1874

*pathmgr\_unlink()* 1876

paths

base name 199

directory name 385

names

maximum length 674, 1867

patterns, matching 647, 943,  
947

truncating 674, 1867

resolving 2341

resource managers

attaching to 2394

detaching from 2413

getting 2453

**\_PATH\_UTMP** 494

pattern matching *See* regular  
expressions

*pause()* 1878

PC Card server

arming 1881

attaching 1884

card insertion/removal,

notification of 1881

CIS (Card Information

Structure), reading 1894

detaching 1886

locking 1891

socket setup information 1888

unlocking 1895

*pccard\_arm()* 1881

**\_PCCARD\_ARM\_INSERT\_REMOVE** 1880

- pccard\_attach()* 1884
- pccard\_detach()* 1886
- `_PCCARD_DEV_AIMS` 1880
- `_PCCARD_DEV_ALL` 1880
- `_PCCARD_DEV_FIXED_DISK` 1880
- `_PCCARD_DEV_GPIB` 1880
- `_PCCARD_DEV_MEMORY` 1880
- `_PCCARD_DEV_NETWORK` 1880
- `_PCCARD_DEV_PARALLEL` 1880
- `_PCCARD_DEV_SCSI` 1880
- `_PCCARD_DEV_SERIAL` 1880
- `_PCCARD_DEV_SOUND` 1880
- `_PCCARD_DEV_VIDEO` 1880
- pccard\_info()* 1888
- pccard\_lock()* 1891
- `_PCCARD_MEMTYPE_ATTRIBUTE` 1893
- `_PCCARD_MEMTYPE_COMMON` 1893
- pccard\_raw\_read()* 1894
- pccard\_unlock()* 1895
- `_PC_CHOWN_RESTRICTED` 674, 1867
- PCI
  - addresses
    - converting 1904
    - testing 1904
  - BIOS, determining if present 1922
  - classes, finding 1912
  - devices
    - attaching 1899
    - configuration, reading 1924, 1926, 1928, 1930
    - configuration, writing 1934, 1937, 1939, 1941
    - detaching 1910
    - finding 1912, 1914
    - rescanning for 1932
    - functions, finding 1912
    - interrupts
      - mapping 1919
      - routing information 1916
    - memory, sharing 1587
    - server
      - attaching 1897
      - detaching 1908
- `pci_dev_info` 1900
- pci\_attach()* 1897
- pci\_attach\_device()* 1899
- `PCIBAD_REGISTER_NUMBER` 1925, 1927, 1929, 1931, 1938, 1940, 1942
- `PCIBUFFER_TOO_SMALL` 1925, 1927, 1929, 1938, 1940
- pci\_detach\_device()* 1910
- `PCI_DEVICE_NOT_FOUND` 1910, 1912, 1913, 1915, 1925
- pci\_find\_class()* 1912
- pci\_find\_device()* 1914
- `PCLINIT_ALL` 1904
- `PCLINIT_BASE0...`
  - `PCLINIT_BASE5` 1904
- `PCLINIT_IRQ` 1903
- `PCLINIT_ROM` 1903
- PCI\_IO\_ADDR()* 1904
- pci\_irq\_routing\_options()* 1916
- PCI\_IS\_IO()* 1904
- PCI\_IS\_MEM()* 1904
- pci\_map\_irq()* 1919
- PCI\_MEM\_ADDR()* 1904
- `PCI_PERSIST` 1903, 1910
- pci\_present()* 1922
- pci\_read\_config()* 1924
- pci\_read\_config16()* 1928
- pci\_read\_config32()* 1930

- pci\_read\_config8()* 1926
- pci\_rescan\_bus()* 1932
- PCI\_ROM\_ADDR()* 1904
- PCLSEARCH\_BUSDEV 1903
- PCLSEARCH\_CLASS 1903
- PCLSEARCH\_VEND 1903
- PCLSEARCH\_VENDEV 1903
- PCLSET\_FAILED 1919
- PCLSHARE 1903
- PCLSUCCESS 1913, 1915, 1922, 1927, 1929, 1931, 1938, 1940, 1942
- PCLUNSUPPORTED\_FUNCTION 1920
- pci\_write\_config()* 1934
- pci\_write\_config16()* 1939
- pci\_write\_config32()* 1941
- pci\_write\_config8()* 1937
- \_PC\_LINK\_MAX 673, 1866
- pclose()* 1943, 1956
- \_PC\_MAX\_CANON 673, 1866
- \_PC\_MAX\_INPUT 673, 1866
- \_PC\_NAME\_MAX 673, 1866
- \_PC\_NO\_TRUNC 674, 1867
- \_PC\_PATH\_MAX 674, 1867
- \_PC\_PIPE\_BUF 674, 1867
- \_PC\_VDISABLE 674, 1867, 3214
- PDU (Protocol Data Unit) *See* SNMP
- peers, getting names of
  - connected 850
- Peripheral Component Interconnect
  - See* PCI
- permissions
  - changing 261, 578
  - files, on creation 3353
  - daemons 1978
- perror()* 1945
- PF\_INET 3141
- PF\_KEY 1320
- PF\_KEY\_V2 1320
- PF\_ROUTE 1052
- pipe()* 1947
- PIPE\_BUF 3551
- pipes
  - bytes, writing atomically 674, 1867
  - closing 1943
  - creating 1947
  - opening 1955
  - reading from 1965, 2306, 2334
- P\_NOWAIT 2933, 2938, 2943, 2947, 2960, 2965, 2969, 2973
- P\_NOWAITO 2933, 2938, 2943, 2947, 2960, 2965, 2969, 2973
- pointers, size of `void` 111
- POOL\_FLAG\_EXIT\_SELF 3218
- POOL\_FLAG\_USE\_SELF 3218
- popen()* 1955
- portable code 103
- ports
  - managing 2494
  - privileged
    - socket, binding to 209
    - socket, getting for 2490
  - reading from 1001, 1003, 1005, 1007, 1009, 1011
  - serial
    - opening 1597
    - reading 1602
    - script, running on 1605
    - writing 1612
  - services, finding for 893

- writing to 1854, 1856, 1858, 1860, 1862, 1864
- POSIX *See also* message queues; semaphores; threads
  - signals 2796
  - standards 103
  - version supported 3137
- \_POSIX\_AIO\_MAX 1442
- \_POSIX\_CHOWN\_RESTRICTED 265, 581, 1423
- \_POSIX\_LOGIN\_NAME\_MAX 829
- posix\_memalign()* 1961
- posix\_mem\_offset()*, *posix\_mem\_offset64()* 1959
- \_POSIX\_THREAD\_SAFE\_FUNCTIONS 788, 829
- P\_OVERLAY 49, 2933, 2938, 2943, 2947, 2960, 2965, 2969, 2973
- pow()*, *powf()* 1963
- PowerPC platforms, variable-length argument lists on 3401
- powers 1963
- P\_PGID 3466
- P\_PID 3466
- pread()*, *pread64()* 1965
- precision, floating-point 667
- printable, testing a character
  - for 1362, 1370, 1389, 1393
- printf()* 1968
- priorities
  - adjusting 1823, 2545
  - getting 860, 2542
  - maximum 2547
  - minimum 2549
  - setting 2561, 2565, 2709
- process groups
  - changing 1991
  - creating 2705, 2727, 2926
  - devices 3178, 3197
  - ID, getting 852, 854
  - joining 2705
  - membership, inheriting 518, 524, 531, 542, 548, 554, 2929
  - pulses, sending 1719
  - remote node 2930
  - session of a controlling terminal 3180
  - setting 2708, 2926
  - signals, sending 1416, 1419, 2767, 2801, 2804
    - SIGHUP 564
    - SIGURG 2909
  - status of 3462, 3469
  - waiting for 3466
- processes *See also* threads
  - address space
    - device I/O memory, mapping 1594
    - limits 2720
    - locking 1582
    - unlocking 1760
  - alarms, scheduling 141, 3345
  - analyzing 1093, 3331
  - arguments
    - auxiliary 198
    - maximum length 3136
    - number of 150
    - parsing 842
    - vector of 151
  - background 360, 1978

- termination, notification
  - of 1980
- child
  - state change, waiting
    - for 3456, 3460, 3463, 3467, 3470
  - zombie, preventing from becoming 2766, 2799
- configurable limits 3136
- connections
  - client, information
    - about 331
    - detaching 251, 335
    - flags, modifying 337
    - information about 340
- controlling terminal, path
  - name 355
- CPU time, maximum 2720
- creating 516, 522, 529, 536, 540, 546, 552, 559, 655, 659, 2927, 2934, 2939, 2944, 2948, 2954, 2961, 2966, 2970, 2974
- data segment, changing space
  - allocated for 211
- debugging 363
- dynamically linked libraries
  - addresses, translating 411
  - closing 413
  - debugging 422
  - errors 415
  - opening 417
  - symbol, getting address
    - of 424
- environment 495
  - clearing 274
  - restoring 1473, 2787
  - saving 2691, 2832
- environment variables
  - defining 2255, 2669
  - deleting 2255, 2669, 3383
  - getting 777
- executable file
  - base name 1993
  - file descriptor 315
  - full path 316
- executing 516, 522, 529, 536, 540, 546, 552, 559
- execution time 300
- execution time limit,
  - getting 2559
- file-mode creation mask 3353
  - daemons 1978
- files, maximum per 3137
- forking 655, 659
- group ID
  - effective 775, 2666, 2675, 2714
  - real 781, 2675, 2714
  - saved 2675
  - supplementary 798, 2680
- guardian, specifying 1987
- I/O priority, requesting 1069, 1078, 1083, 1085, 1087, 1095, 1097, 1100, 1102, 3246
- ID, getting 856, 858
- interrupts
  - disabling 1085
  - enabling 1087
  - events 1077, 1083
  - handlers 1069, 1083, 1095, 1100
  - handlers, idle 1089

- masking 1097, 1102
- waiting for 1104
- maximum per real user
  - ID 3136
- memory, sharing 1587
- message channels, attaching to 327
- name 316, 1993
- parent
  - blocking 3416
  - ID, getting 858
- priority
  - adjusting 1823, 2545
  - getting 860, 2542
  - maximum 2547
  - minimum 2549
  - setting 2561, 2565, 2709
- processor affinity 3246
- program entry function 1495
- scheduling policy
  - getting 2551
  - setting 2565
- sessions 897, 2708, 2727
- set-group ID 3137
- set-user ID 3137
- SIGALRM sending to 3260, 3285, 3345
- signals
  - actions for 2764, 2791, 2796
  - information about 2845
  - pending 2824
  - queuing 2829
  - raising 2283
  - sending 1416, 1419, 2804
  - suspending until
    - delivered 2813
    - waiting for 1878, 2816, 2838, 2843, 2845
- spawning 2927, 2934, 2939, 2944, 2948, 2954, 2961, 2966, 2970, 2974
- spawning and blocking 3416
- supplementary group IDs,
  - maximum 3136
- suspending 368, 1782, 1796, 1797, 2851, 3385, 3456, 3460, 3463, 3467, 3470
- system commands,
  - executing 3158
- system-wide events
  - notification of 1981
  - triggering 1985
- terminating 113, 563, 566
  - diagnostics 158
  - functions to be called,
    - registering 167
- time
  - clock ID 286
  - clock ticks 279
- time-accounting
  - information 3306
- user ID
  - effective 779, 2672, 2717, 2737
  - real 925, 2717, 2737
  - saved 2737
  - user name, getting 827, 829
  - yielding 2567
- zombies, preventing children
  - from becoming 2766, 2799
- processor affinity 3246
- Processor Version Register 3152

- procmgr\_daemon()* 1978
- PROCMGR\_DAEMON\_KEEPUKMASK 1978
- PROCMGR\_DAEMON\_NOCHDIR 1978
- PROCMGR\_DAEMON\_NOCLOSE 1978
- PROCMGR\_DAEMON\_NODEVNULL 1978
- PROCMGR\_EVENT\_DAEMON\_DEATH 1980
- procmgr\_event\_notify()* 1981
- PROCMGR\_EVENT\_SYNC 1980
- procmgr\_event\_trigger()* 1985
- procmgr\_guardian()* 1987
- procmgr\_session()* 1991
- PROCMGR\_SESSION\_SETPGRP 1991
- PROCMGR\_SESSION\_SETSID 1991
- PROCMGR\_SESSION\_SIGNAL\_LEADER 1991
- PROCMGR\_SESSION\_SIGNAL\_PGRP 1991
- PROCMGR\_SESSION\_SIGNAL\_PID 1991
- PROCMGR\_SESSION\_TCSETSID 1991
- \_\_progrname* 1993
- program entry function 1495
- PROT\_EXEC 1207, 1584, 1593, 1625
- PROT\_NOCACHE 1207, 1584, 1593, 1625
- PROT\_NONE 1207, 1584, 1593, 1625
- Protocol Data Unit (PDU) *See* SNMP
- protocols
  - database
    - closing 490
    - entries, getting 862, 864, 866
    - entry structure 1994
    - opening 2711
  - ICMP (Internet Control Message Protocol) 983
  - ICMP6 (Internet Control Message Protocol v6) 985
  - INET6 (Internet protocol v6 family) 1051
  - interprocess
    - communication 3377
  - IP (Internet Protocol) 1313
  - IP6 (Internet Protocol v6) 1338
  - IPsec (secure Internet Protocol) 1320
  - TCP (Transmission Control Protocol) 3189
  - UDP (User Datagram Protocol) 3348
- protoent** 1994
- PROT\_READ 1207, 1584, 1593, 1625
- PROT\_WRITE 1207, 1584, 1593, 1625
- proxy server (SOCKS) 3571
- pseudo-random numbers
  - double** 431, 498
  - int** 2286, 2288
  - long**
    - nonnegative 1476, 1825
    - signed 1414, 1656, 2290
  - seed, setting 2605, 2982, 2984, 2986
  - sequence, initializing 2984
  - state
    - initializing 1063, 1426
    - switching 2733
- pseudo-ttys

- opening 659, 1852
- preparing for a login 659, 1471
- pthread\_attr\_t** 2089
- PTHREAD\_COND\_INITIALIZER**
  - 2066
- pthread\_abort()* 1995
- PTHREAD\_ABORTED 1995
- pthread\_atfork()* 1997
- pthread\_attr\_destroy()* 1999
- pthread\_attr\_getdetachstate()* 2001
- pthread\_attr\_getguardsize()* 2003
- pthread\_attr\_getinheritsched()*
  - 2005
- pthread\_attr\_getschedparam()* 2007
- pthread\_attr\_getschedpolicy()*
  - 2009
- pthread\_attr\_getscope()* 2011
- pthread\_attr\_getstackaddr()* 2013
- pthread\_attr\_getstacklazy()* 2015
- pthread\_attr\_getstacksize()* 2017
- pthread\_attr\_init()* 2019
- pthread\_attr\_setdetachstate()* 2022
- pthread\_attr\_setguardsize()* 2024
- pthread\_attr\_setinheritsched()*
  - 2027
- pthread\_attr\_setschedparam()* 2029
- pthread\_attr\_setschedpolicy()* 2031
- pthread\_attr\_setscope()* 2033
- pthread\_attr\_setstackaddr()* 2035
- pthread\_attr\_setstacklazy()* 2037
- pthread\_attr\_setstacksize()* 2039
- pthread\_barrierattr\_destroy()* 2047
- pthread\_barrierattr\_getpshared()*
  - 2049
- pthread\_barrierattr\_init()* 2051
- pthread\_barrierattr\_setpshared()*
  - 2053
- pthread\_barrier\_destroy()* 2041
- pthread\_barrier\_init()* 2043
- PTHREAD\_BARRIER\_INITIALIZER()**
  - 2043
- pthread\_barrier\_wait()* 2045
- PTHREAD\_CANCEL 3234
- pthread\_cancel()* 2055
- PTHREAD\_CANCEL\_ASYNCHRONOUS 2091,
  - 2198, 3241
- PTHREAD\_CANCEL\_DEFERRED 2091,
  - 2198, 3241, 3242
- PTHREAD\_CANCEL\_DISABLE 2091,
  - 2196
- PTHREAD\_CANCELED 2108, 2236
- PTHREAD\_CANCEL\_ENABLE 2091,
  - 2196, 3242
- pthread\_cleanup\_pop()* 2057
- pthread\_cleanup\_push()* 2059
- pthread\_condattr\_destroy()* 2077
- pthread\_condattr\_getclock()* 2079
- pthread\_condattr\_getpshared()*
  - 2081
- pthread\_condattr\_init()* 2083
- pthread\_condattr\_setclock()* 2085
- pthread\_condattr\_setpshared()*
  - 2087
- pthread\_cond\_broadcast()* 2062
- pthread\_cond\_destroy()* 2064
- pthread\_cond\_init()* 2066
- PTHREAD\_COND\_INITIALIZER 3104,
  - 3107
- pthread\_cond\_signal()* 2068
- pthread\_cond\_timedwait()* 2070
- pthread\_cond\_wait()* 2074
- pthread\_create()* 2090
- PTHREAD\_CREATE\_DETACHED 2022,
  - 3241



- PTHREAD\_CREATE\_JOINABLE 2019,  
 2022, 3241  
 PTHREAD\_DESTRUCTOR\_ITERATIONS  
 2111  
*pthread\_detach()* 2094  
*pthread\_equal()* 2096  
*pthread\_exit()* 2098  
 PTHREAD\_EXPLICIT\_SCHED 2027,  
 2029, 2031, 3240, 3241  
*pthread\_getconcurrency()* 2100  
*pthread\_getcpuclockid()* 2102  
*pthread\_getschedparam()* 2104  
*pthread\_getspecific()* 2106  
 PTHREAD\_INHERIT\_SCHED 2019,  
 2027  
*pthread\_join()* 2108  
*pthread\_key\_create()* 2110  
*pthread\_key\_delete()* 2114  
*pthread\_kill()* 2116  
 PTHREAD\_MULTISIG\_ALLOW 2091,  
 3241  
 PTHREAD\_MULTISIG\_DISALLOW 2091,  
 3241  
*pthread\_mutexattr\_destroy()* 2137  
*pthread\_mutexattr\_getprioceiling()*  
 2139  
*pthread\_mutexattr\_getprotocol()*  
 2141  
*pthread\_mutexattr\_getpshared()*  
 2143  
*pthread\_mutexattr\_getrecursive()*  
 2145  
*pthread\_mutexattr\_gettype()* 2147  
*pthread\_mutexattr\_init()* 2150  
*pthread\_mutexattr\_setprioceiling()*  
 2152  
*pthread\_mutexattr\_setprotocol()*  
 2154  
*pthread\_mutexattr\_setpshared()*  
 2156  
*pthread\_mutexattr\_setrecursive()*  
 2158  
*pthread\_mutexattr\_settype()* 2161  
 PTHREAD\_MUTEX\_DEFAULT 2161  
*pthread\_mutex\_destroy()* 2118  
 PTHREAD\_MUTEX\_ERRORCHECK 2160  
*pthread\_mutex\_getprioceiling()*  
 2120  
*pthread\_mutex\_init()* 2122  
 PTHREAD\_MUTEX\_INITIALIZER 2122  
*pthread\_mutex\_lock()* 2124  
 PTHREAD\_MUTEX\_NORMAL 2160  
 PTHREAD\_MUTEX\_RECURSIVE 2160  
*pthread\_mutex\_setprioceiling()*  
 2128  
*pthread\_mutex\_timedlock()* 2130  
*pthread\_mutex\_trylock()* 2133  
*pthread\_mutex\_unlock()* 2135  
*pthread\_once()* 2163  
 PTHREAD\_ONCE\_INIT 2163  
 PTHREAD\_PRIO\_INHERIT 2150,  
 2154, 3132  
 PTHREAD\_PRIO\_NONE 2154  
 PTHREAD\_PRIO\_PROTECT 2154,  
 3132  
 PTHREAD\_PROCESS\_PRIVATE 2053,  
 2083, 2156, 2193, 2226  
 PTHREAD\_PROCESS\_SHARED 2053,  
 2083, 2087, 2156, 2193,  
 2226  
 PTHREAD\_RECURSIVE\_DISABLE 2145,  
 2150, 2158

- PTHREAD\_RECURSIVE\_ENABLE 2145,  
2158
- PTHREAD\_RMUTEX\_INITIALIZER 2122
- pthread\_rwlockattr\_destroy()* 2187
- pthread\_rwlockattr\_getpshared()*  
2189
- pthread\_rwlockattr\_init()* 2191
- pthread\_rwlockattr\_setpshared()*  
2193
- pthread\_rwlock\_destroy()* 2166
- pthread\_rwlock\_init()* 2168
- PTHREAD\_RWLOCK\_INITIALIZER 2168
- pthread\_rwlock\_rdlock()* 2171
- pthread\_rwlock\_timedrdlock()* 2173
- pthread\_rwlock\_timedwrlock()* 2176
- pthread\_rwlock\_tryrdlock()* 2179
- pthread\_rwlock\_trywrlock()* 2181
- pthread\_rwlock\_unlock()* 2183
- pthread\_rwlock\_wrlock()* 2185
- PTHREAD\_SCOPE\_PROCESS 3241
- PTHREAD\_SCOPE\_SYSTEM 2019,  
2033, 3241
- pthread\_self()* 2195
- pthread\_setcancelstate()* 2196
- pthread\_setcanceltype()* 2198
- pthread\_setconcurrency()* 2200
- pthread\_setschedparam()* 2203
- pthread\_setspecific()* 2204
- pthread\_sigmask()* 2206
- pthread\_sleepon\_broadcast()* 2208
- pthread\_sleepon\_lock()* 2210
- pthread\_sleepon\_signal()* 2212
- pthread\_sleepon\_timedwait()* 2214
- pthread\_sleepon\_unlock()* 2218
- pthread\_sleepon\_wait()* 2220
- pthread\_spin\_destroy()* 2224
- pthread\_spin\_init()* 2226
- pthread\_spin\_lock()* 2228
- pthread\_spin\_trylock()* 2230
- pthread\_spin\_unlock()* 2232
- PTHREAD\_STACK\_LAZY 2037
- PTHREAD\_STACK\_MIN 2035,  
2036, 2039, 2040, 3239,  
3244
- PTHREAD\_STACK\_NOTLAZY 2037
- pthread\_testcancel()* 2234
- pthread\_timedjoin()* 2235
- \_PTR\_BITS\_ 111
- pulse\_attach()* 2241
- \_PULSE\_CODE\_COIDDEATH 250,  
2238
- \_PULSE\_CODE\_DISCONNECT 251,  
2238
- \_PULSE\_CODE\_MAXAVAIL 1718,  
2238, 2779
- \_PULSE\_CODE\_MINAVAIL 1718,  
2238, 2779
- \_PULSE\_CODE\_NET\_ACK 2238
- \_PULSE\_CODE\_NET\_DETACH 2238
- \_PULSE\_CODE\_NET\_UNBLOCK 2238
- \_PULSE\_CODE\_THREADDEATH 252,  
402, 2238
- \_PULSE\_CODE\_UNBLOCK 252,  
2238
- pulse\_detach()* 2244
- pulses
- compression 1719
  - dispatch interface
    - attaching 2241
    - detaching 2244
  - queueing 1719
  - receiving 1694, 1697
  - sending 1719
  - structure 2238

\_PULSE\_SUBTYPE 2238  
 \_PULSE\_TYPE 2238  
 punctuation, testing a character  
     for 1372, 1395  
*putc()* 2247  
*putchar()* 2251  
*putchar\_unlocked()* 2253  
*putc\_unlocked()* 2249  
*putenv()* 537, 560, 2255, 2934,  
     2944, 2961, 2970  
*puts()* 2258  
*putspent()* 2260  
*pututline()* 2263  
*putw()* 2266  
*putwc()* 2268, 2270  
 P\_WAIT 49, 2933, 2938, 2943,  
     2947, 2960, 2965, 2969,  
     2973  
*pwrite()* 2272

## Q

\_\_QNX\_\_ 111  
 QNX 4 classification 106  
 QNX Neutrino classification 106  
*qnx\_crypt()* 2275  
 \_\_QNXNTO\_\_ 111  
 QoS (Quality of Service) 1807  
*qsort()* 2277  
**qtime** 3154  
 Quality of Service (QoS) 1807  
 QUERY 2378  
 quick sort 2277  
 quotients  
     integer 408

long integer 1430

## R

*Raccept()* 2281  
 radix-independent exponents 1468,  
     2523, 2526  
*raise()* 1473, 2283  
*rand()* 2286  
 RAND\_MAX 2286, 2288  
 random numbers  
     **double** 431, 498  
     **int** 2286, 2288  
     **long**  
         nonnegative 1476, 1825  
         signed 1414, 1656, 2290  
     seed, setting 2605, 2982, 2984,  
         2986  
     sequence, initializing 2984  
     state  
         initializing 1063, 1426  
         switching 2733  
*random()* 2290  
*rand\_r()* 2288  
 raw input mode  
     buffer 673, 1866  
     conditions for input  
         request 2319  
     **FORWARD** qualifier 2320  
     **MIN** qualifier 2319  
     **TIME** qualifier 2319  
     **TIMEOUT** qualifier 2320  
*Rbind()* 2293  
*rcmd()* 2295, 3572  
*Rconnect()* 2298

- rcvid* (receive identifier) 1690, 1701
  - checking validity of 1746
- rdchk()* 2300
- read()* 1135, 2306, 2318
  - resource managers, implementing in 1254, 1256
- read-write locks
  - attributes
    - creating 2191
    - destroying 2187
    - process-shared 2189, 2193
  - destroying 2166, 3114
  - initializing 2168, 3133
  - locking
    - for reading 2171, 2173, 2179
    - for writing 2176, 2181, 2185
  - unlocking 2183
- readblock()* 2315
- readcond()* 2318
- readdir()* 1843, 2324
- readdir\_r()* 2328
- readlink()* 2331
  - resource managers, implementing in 1260
- read\_main\_config\_file()* 2311
- readv()* 2334
- realloc()* 2338
- realpath()* 2341
- realtime timers
  - busy-waiting 1784, 1789, 1791, 1793
  - calibrating 1786
  - creating 3259
  - destroying 3263
  - expiry status 3265
  - overruns 3267
  - setting 3295
  - threads 294
  - time
    - getting 3269
    - setting 3272
- rebooting 3150
- receive identifier *See rcvid*
- RECEIVED\_MESSAGE 2899
- receiving
  - messages 1689, 1700
    - from a socket 2344, 2347, 2351
  - pulses 1694, 1697
- re\_comp()* 2302
- recv()* 2344
- recvfrom()* 2347
- recvmsg()* 2351
- re\_exec()* 2304
- regcomp()* 2354
- regerror()* 2359
- regex\_t** 2354
- regexexec()* 2362
- REG\_EXTENDED 2354
- regfree()* 2364
- REG\_ICASE 2354
- registers
  - devices
    - access to, gaining and relinquishing 1591, 1764
    - reading 1924, 1926, 1928, 1930
    - writing 1934, 1937, 1939, 1941
  - floating-point exceptions 664
  - Machine Status Register 3152

- Processor Version
  - Register 3152
- real-mode software
  - interrupts 1107
- TSC (Time Stamp Counter) 284
- regmatch\_t** 2362
- REG\_NEWLINE 2354
- REG\_NOSUB 2354, 2362
- REG\_NOTBOL 2361
- REG\_NOTEOL 2361
- regular expressions
  - basic 2355
  - compiling 2302, 2354
  - errors, explaining 2359
  - extended 2356
  - freeing 2364
  - string, comparing to 2304, 2362
- remainder()*, *remainderf()* 2366
- remainders
  - floating point 433, 2366
  - integer 408
  - long integer 1430
- remote hosts
  - commands, executing on 2295, 2488
  - identity, checking 2518
- remove()* 2368
- rename()* 2371
- RES\_DEBUG 2374
- RES\_DEFNAMES 2375, 2387
- RES\_DNSRCH 2375, 2387
- residue, floating point 643
- RES\_INIT 2375
- res\_init()* 2374
- resmgr\_attr\_t** 2395
- resmgr\_connect\_funcs\_t** 2404
- resmgr\_context\_t** 2411
- resmgr\_io\_funcs\_t** 2431
- resmgr\_attach()* 2394
- resmgr\_block()* 2401
- \_RESMGR\_CONNECT\_NFUNCS 1179
- resmgr\_context\_alloc()* 2406
- resmgr\_context\_free()* 2409
- \_RESMGR\_DEFAULT 1285
- resmgr\_detach()* 2413
- \_RESMGR\_DETACH\_ALL 2413
- \_RESMGR\_DETACH\_PATHNAME 2413
- resmgr\_devino()* 2417
- RESMGR\_FLAG\_AFTER 1618, 1620
- \_RESMGR\_FLAG\_AFTER 2396
- RESMGR\_FLAG\_ATTACH\_OTHERFUNC 2396
- RESMGR\_FLAG\_BEFORE 1617, 1620
- \_RESMGR\_FLAG\_BEFORE 2396
- \_RESMGR\_FLAG\_DIR 2396
- \_RESMGR\_FLAG\_FTYPEONLY 2397
- \_RESMGR\_FLAG\_OPAQUE 2397
- RESMGR\_FLAG\_OPAQUE 1618, 1620
- \_RESMGR\_FLAG\_SELF 2397
- resmgr\_handle\_grow()* 2423
- resmgr\_handler()* 2425
- resmgr\_io\_func()* 2428
- resmgr\_iofuncs()* 2436
- \_RESMGR\_IO\_NFUNCS 1179
- resmgr\_msgread()* 2438
- resmgr\_msgreadv()* 2440
- resmgr\_msgwrite()* 2442
- resmgr\_msgwritev()* 2444

- `_RESMGR_NOREPLY` 1285
- `_RESMGR_NPARTS()` 2446
- `_resmgr_ocb()` 2448
- `resmgr_open_bind()` 2450
- `resmgr_pathname()` 2453
- `_RESMGR_PATHNAME_LOCALPATH` 2453
- `_RESMGR_PTR()` 2456
- `_RESMGR_STATUS()` 2458
- `resmgr_unbind()` 2460
- `res_mkquery()` 2378
- `resolv.conf`, contents of 319
- resolver routines
  - errors 953, 958, 967
  - initializing 2374
  - Internet domain names
    - compressing 427
    - expanding 429
  - options 2374
  - queries 2378, 2381, 2384, 2387, 2389
- resource database manager
  - about 2494
  - device numbers
    - attaching 2510
    - detaching 2513
  - resources
    - creating 2501
    - destroying 2505
    - querying 2516
    - reserving 2494
    - returning 2507
- resource managers
  - access, checking 1141
  - arming for notification 1307
  - attaching 2394
  - attributes
    - initializing 1127
    - locking 1129, 1134, 1137, 1195
  - structure 1132
  - time members,
    - updating 1283
  - unlocking 1139
- clients
  - information about 1155
  - unblocking 1285, 1287
- connect functions 2404
  - default values, setting 1179
  - open 1241
- connect messages
  - file type reply 1118
  - link reply 1120
  - structure 1111
- connection IDs 327
- context
  - allocating 2406
  - freeing 2409
  - structure 2411
- database, expanding
  - capacity 2423
- detaching 2413
- device number, getting 2417
- device-control commands 372
- function tables,
  - initializing 1179
- helper functions
  - `chmod` 1145
  - `chown` 1150
  - `close` 1157
  - `devctl` 1166
  - `fdinfo` 1173
  - `link` 1182
  - `lseek` 1197
  - `mknod` 1202

- mmap* 1205
- open* 1236
- openfd* 1243
- pathconf* 1249
- read* 1256
- readlink* 1260
- rename* 1264
- space* 1267
- stat* 1271
- sync* 1280
- unlink* 1290
- utime* 1295
- write* 1303
- I/O functions
  - chmod* 1148
  - chown* 1153
  - client connection, getting
    - for 2436
  - close* 1160
  - default values, setting 1179
  - devctl* 1170
  - fdinfo* 1175
  - lock* 1190
  - lseek* 1200
  - mmap* 1209
  - OCB, close 1164
  - OCB, lock 1195
  - OCB, unlock 1293
  - openfd* 1247
  - pathconf* 1252
  - read* 1254
  - retrieving 2428
  - stat* 1273
  - structure 2431
  - sync* 1278
  - unblock 1287
  - utime* 1298
  - write* 1301
- inode number, getting 2417
- iov\_t**
  - filling 2456
  - getting 2446
- locks (not implemented) 1186, 1188, 1193
- messages
  - blocking while waiting for 2401
  - handling 2425
  - reading 2438, 2440
  - writing 2442, 2444
- notification
  - arming for 1307
  - installing, polling, and removing 1213
  - removing for a client 1218
  - triggering 1220
- Open Control Block (OCB)
  - allocating 1225
  - attaching 1224, 2450
  - detaching 1162, 1228, 2460
  - freeing 1231
  - getting 2448
  - structure 1234
- path
  - attaching to 2394
  - detaching from 2413
  - getting 2453
- server attributes, getting 1125
- status, returning 2458
- synchronization, checking to see if required 1276
- threads in 1134
- resources, system
  - creating 2501

- destroying 2505
- limits
  - getting 773, 881
  - setting 2719
- querying 2516
- reserving 2494
- returning 2507
- usage, getting 884
- res\_query()* 2381
- res\_querydomain()* 2384
- RES\_RECURSE 2375
- res\_search()* 2387
- res\_send()* 2389
- RES\_STAYOPEN 2375
- RES\_USEVC 2375
- rewind()* 2462
- rewinddir()* 2465
- rftp** 3571
- Rgetsockname()* 2468
- rindex()* 2470
- rint()*, *rintf()* 2472
- RLIM\_INFINITY 2720, 2723
- RLIMIT\_AS 2720
- RLIMIT\_CORE 2720
- RLIMIT\_CPU 2720
- RLIMIT\_DATA 2720
- RLIMIT\_FSIZE 2720, 2723
- RLIMIT\_NOFILE 773, 2720, 2723
- RLIMIT\_STACK 2720
- RLIMIT\_VMEM 2720
- RLIM\_SAVED\_CUR 2723
- RLIM\_SAVED\_MAX 2723
- Rlisten()* 2475
- rmdir()* 2477
- R\_OK 120, 456
- root directory, changing 268
- roots
  - cube 226
  - square 2980
- round-robin scheduling 2031
- rounding
  - floating point 2472
  - integers 228, 639
  - mode, floating-point 670
- ROUTE (system packet forwarding database) 2480
- Rrcmd()* 2488
- rresvport()* 2490
- Rselect()* 2493
- rshd** 2295
- rsrc\_alloc\_t** 2501
- rsrc\_request\_t** 2495
- rsrcdbmgr\_attach()* 2494
- rsrcdbmgr\_create()* 2501
- rsrcdbmgr\_destroy()* 2505
- rsrcdbmgr\_detach()* 2507
- rsrcdbmgr\_devno\_attach()* 2510
- rsrcdbmgr\_devno\_detach()* 2513
- RSRCDBMGR\_DMA\_CHANNEL 2496, 2502, 2515
- RSRCDBMGR\_FLAG\_ALIGN 2496
- RSRCDBMGR\_FLAG\_RANGE 2496
- RSRCDBMGR\_FLAG\_RSVP 2502
- RSRCDBMGR\_FLAG\_SHARE 2496
- RSRCDBMGR\_FLAG\_TOPDOWN 2496
- RSRCDBMGR\_IO\_PORT 2496, 2502, 2515
- RSRCDBMGR\_IRQ 2496, 2502, 2515
- RSRCDBMGR\_MEMORY 2496, 2502, 2515
- RSRCDBMGR\_PCI\_MEMORY 2496, 2502, 2515
- rsrcdbmgr\_query()* 2516



RSRCMGR\_IRQ 2502, 2515  
**rt\_metrics** 2483  
**rt\_msghdr** 2483  
RTA\_AUTHOR 2484  
RTA\_BRD 2484  
RTA\_DST 2484  
RTA\_GATEWAY 2484  
RTA\_GENMASK 2484  
RTA\_IFA 2484  
RTA\_IFP 2484  
RTA\_NETMASK 2484  
**rtelnet** 3571  
RTF\_BLACKHOLE 2484  
RTF\_CLONING 2484  
RTF\_DONE 2484  
RTF\_DYNAMIC 2484  
RTF\_GATEWAY 2484  
RTF\_HOST 2484  
RTF\_LLINFO 2484  
RTF\_MASK 2484  
RTF\_MODIFIED 2484  
RTF\_PROTO1 2484  
RTF\_PROTO2 2484  
RTF\_REJECT 2484  
RTF\_STATIC 2484  
RTF\_UP 2484  
RTF\_XRESOLVE 2484  
RTLDD\_DEFAULT 424  
RTLDD\_GLOBAL 420  
RTLDD\_GROUP 421  
RTLDD\_LAZY 420  
RTLDD\_LOCAL 420  
RTLDD\_NOW 420  
RTLDD\_WORLD 421  
RTM\_ADD 2482  
RTM\_CHANGE 2482  
RTM\_DELADDR 2482, 2483

RTM\_DELETE 2482  
RTM\_GET 2482  
RTM\_IFINFO 2482, 2483  
RTM\_LOSING 2482  
RTM\_MISS 2482  
RTM\_NEWADDR 2482, 2483  
RTM\_REDIRECT 2482  
RTM\_RESOLVE 2482  
RTV\_EXPIRE 2484  
RTV\_HOPCOUNT 2484  
RTV\_MTU 2484  
RTV\_RPIPE 2484  
RTV\_RTT 2484  
RTV\_RTTVAR 2484  
RTV\_SPIPE 2484  
RTV\_SSTHRESH 2484  
RUN\_LVL 929, 3393  
**rusage** 884  
*ruserok()* 2518

## S

SA\_NOCLDSTOP 2792, 2798  
SA\_ONSTACK 517  
SA\_SIGINFO 1443, 2765, 2799,  
2829  
SAT (System Analysis  
Toolkit) 1093, 3331  
*sbrk()* 2520  
*scalb(), scalbf()* 2523  
*scalbn(), scalbnf()* 2526  
*\_scalloc()* 2529  
*scandir()* 2531  
*scanf()* 2533  
\_SC\_ARG\_MAX 3136

- .\_SC\_CHILD\_MAX 3136
- .\_SC\_CLK\_TCK 3136
- .\_SC\_GETGR\_R\_SIZE\_MAX 788
- .\_SC\_GETPW\_R\_SIZE\_MAX 873, 878, 911, 915
- sched\_param** 2553
- SCHED\_FIFO 2031, 2547, 2549, 2564, 2570, 2573
- SchedGet()*, *SchedGet\_r()* 2570
- sched\_getparam()* 2542
- sched\_get\_priority\_adjust()* 2545
- sched\_get\_priority\_max()* 2547
- sched\_get\_priority\_min()* 2549
- sched\_getscheduler()* 2551
- SchedInfo()*, *SchedInfo\_r()* 2574
- SCHED\_NOCHANGE 2031, 3240
- SCHED\_OTHER 2031, 2547, 2549, 2564, 2570, 2573
- SCHED\_RR 2031, 2547, 2549, 2564, 2570, 2573
- sched\_rr\_get\_interval()* 2559
- SchedSet()*, *SchedSet\_r()* 2577
- sched\_setparam()* 2561
- sched\_setscheduler()* 2565
- SCHED\_SPORADIC 2031, 2570
- scheduling
  - information, getting 2574
  - parameters 2553
  - threads, getting for 2007, 2029, 2104, 2570
  - threads, inheriting 2005, 2027, 3241
  - threads, setting for 2203, 2577
  - policy
    - don't change 2031
    - FIFO 2031
    - other 2031
    - processes, getting for 2551
    - processes, setting for 2565
    - round-robin 2031
    - sporadic 2031
    - threads, getting for 2009, 2104, 2570
    - threads, inheriting 2005, 2027, 3241
    - threads, setting for 2031, 2203, 2577, 3240
    - yielding 2579
- sched\_yield()* 2567
- SchedYield()*, *SchedYield\_r()* 2579
- .\_SC\_JOB\_CONTROL 3137
- SCM\_RIGHTS 3378
- .\_SC\_NGROUPS\_MAX 3136
- scoped addresses 1052
- .\_SC\_OPEN\_MAX 3137
- .\_SC\_PAGESIZE 2024
- scripts, running 536, 559, 2925, 2928, 2934, 2939, 2944, 2948, 2955, 2961, 2966, 2970, 2974
- .\_SC\_SAVED\_IDS 3137
- .\_SC\_VERSION 3137
- searchenv()* 2602
- secure Internet Protocol *See* IPsec
- secure RPC domain 319
- Security Policy Database (SPD) 1324
- seed48()* 2605
- SEEK\_CUR 591, 705, 1197, 1268, 1481, 1491
- seekdir()* 2607
- SEEK\_END 591, 705, 1197, 1269, 1481, 1491

- SEEK\_SET 589, 591, 705, 1197, 1269, 1481, 1491
- segments
  - data
    - maximum size 2720
  - data, end of 459, 462
  - text
    - beginning of 217
    - end of 515
- select\_attr\_t** 2615
- select()* 118, 2610, 3572
  - data from *snmp\_select\_info()* 2892
- select\_attach()* 2616
- select\_detach()* 2619
- SELECT\_FLAG\_EXCEPT 2616
- SELECT\_FLAG\_READ 2616
- SELECT\_FLAG\_REARM 2616
- SELECT\_FLAG\_SRVEXCEPT 2617
- SELECT\_FLAG\_WRITE 2616
- select\_query()* 2623
- semaphores
  - named
    - accessing and creating 2635
    - closing 2625
    - destroying 2646
  - posting 2639, 3127
  - unnamed
    - destroying 2627, 3114
    - initializing 2631, 3133
  - value
    - decrementing 2641, 2644, 2648, 3129
    - getting 2629
    - incrementing 2639, 3127
    - setting 2631, 2635
    - waiting on 2648, 3129
    - with a time limit 2641
    - without blocking 2644
- sem\_close()* 2625
- sem\_destroy()* 2627
- sem\_getvalue()* 2629
- sem\_init()* 2631
- sem\_open()* 2635
- sem\_post()* 2639
- sem\_timedwait()* 2641
- sem\_trywait()* 2644
- sem\_unlink()* 2646
- SEM\_VALUE\_MAX 2631, 2632, 2635
- sem\_wait()* 2648
- send()* 2651
- sendmsg()* 2653
- sendto()* 2657
- serial number, getting 319
- serial ports
  - opening 1597
  - reading 1602
  - script, running on 1605
  - writing 1612
- servent** 2659
- server attributes, getting 1125
- servers
  - connections
    - creating 248
    - destroying 255
    - information about 331, 340
  - data server
    - applications, registering and deregistering 440, 446
    - variables, creating and destroying 435, 438
    - variables, flags 442

- variables, getting and setting 444, 448
- PCI
  - attaching 1897
  - detaching 1908
- services
  - database
    - closing 492
    - entries, getting 891, 893, 895
    - entry structure 2659
    - opening 2725
- sessions
  - character device terminal
    - drivers, support for 1991
  - controlling terminal 3180
  - creating 1471, 2727
  - current 2451
  - disassociating 564
  - ID, getting 897
  - leader, creating 2708
  - membership, inheriting 518, 524, 531, 542, 548, 554, 2955
  - remote node 2930
  - system daemons 1978
    - termination, notification of 1980
- setbuf()* 2660
- setbuffer()* 2662
- setdomainname()* 2664
- setegid()* 2666
- setenv()* 537, 560, 2669, 2934, 2944, 2961, 2970
- seteuid()* 2672
- setgid()* 2675
- setgrent()* 2678
- setgroups()* 2680
- sethostent()* 800, 807, 2682
- sethostname()* 2684
- SETIOV()* 2686
- setitimer()* 2688
- setjmp()* 2691
- setkey()* 2694
- setlinebuf()* 2696
- setlocale()* 2699, 3090
- setlogmask()* 2701
- setnetent()* 2703
- setpgid()* 2705
- setpgrp()* 2708
- setprio()* 2709
- setprotoent()* 2711
- setpwent()* 2713
- setregid()* 2714
- SET\_REQ\_MSG 2887
- setreuid()* 2717
- setrlimit()*, *setrlimit64()* 2719
- setservent()* 2725
- setsid()* 2727
- setsockopt()* 2730
- setspent()* 2732
- setstate()* 2733
- settimeofday()* 2735
- setuid()* 2737
  - not honoring on mounted filesystems 1617
- setutent()* 2740
- setvbuf()* 2742
- \_sfree()* 2745
- sh** 1955
- shadow password database
  - closing 493
  - entries
    - reading 618, 911, 915

- structure 2260
  - writing 2260
- rewinding 2732
- shared locks 634
- shared memory
  - access protection,
    - changing 1625
  - attributes, modifying 2748
  - mapping 1585
  - opening 2754
  - removing 2760
  - unmapping 1762
- shared objects
  - addresses, translating 411
  - closing 413
  - debugging 422
  - errors 415
  - opening 417
  - symbol, getting address of 424
- SH\_COMPAT 1115, 1245, 2918, 2922
- SH\_DENYNO 1115, 1245, 2918, 2922
- SH\_DENYRD 1115, 1245, 2918, 2922
- SH\_DENYRW 1115, 1245, 2918, 2922
- SH\_DENYWR 1115, 1245, 2918, 2922
- SHELL** 274, 1955, 3158
- shell scripts, running 536, 559, 2925, 2928, 2934, 2939, 2944, 2948, 2955, 2961, 2966, 2970, 2974
- shm\_ctl()* 2748
- SHMCTL\_ANON 2747
- SHMCTL\_GLOBAL 2747
- SHMCTL\_LOWERPROT 2747
- SHMCTL\_PHYS 2747
- SHMCTL\_PRIV 2747
- shm\_open()* 2754
- shm\_unlink()* 2760
- shutdown()* 2762
- SLASYNCIO 2767, 2800, 2806
- side channels 327
- S\_IEXEC 1115, 2995
- S\_IFBLK 1114, 1121, 2996
- S\_IFCHR 1114, 1121, 2996
- S\_IFDIR 1114, 1121, 1569, 2996
- S\_IFIFO 1114, 1121, 1569, 2996
- S\_IFLNK 1114, 1121, 2996
- S\_IFMT 1114, 1121, 2996
- S\_IFNAM 1115, 1121, 2996
- S\_IFREG 1115, 1121, 2996
- S\_IFSOCK 1115, 1121, 1360, 2996
- SIGABRT 113, 2796
- sigaction** 2798
- sigaction()* 2764
- sigaddset()* 2770
- SIGALRM 141, 826, 2689, **2797**, 3282
  - process, sending to 3260, 3285, 3345
- SIG\_BLOCK 2826
- sigblock()* 2772
- SIGBUS 1586, 2797
- SIGCHLD 564, 884, 2765, 2792, 2793, **2797**, 3467
  - default actions 2765, 2799
  - ignoring 2766, 2799
- SIGCONT 564, 2772, 2797
  - default actions 2765
- sigdelset()* 2774
- SIG\_DFL 517, 2765, 2791

- sigemptyset()* 2776
- SIGEMT 2796
- SIG\_ERR 2793
- sigevent** 1309, 1661, 2778
- SIGEV\_INTR 1072, 1078, 1104,  
1633, 1662, 2778
- SIGEV\_INTR\_INIT()* 2779
- SIGEV\_NONE 1443, 2778
- SIGEV\_NONE\_INIT()* 2779
- SIGEV\_PULSE 1072, 1078, 1633,  
1662, 2778, 3260
- SIGEV\_PULSE\_INIT()* 2780
- SIGEV\_SIGNAL 1073, 1078, 1443,  
1633, 1662, 2778, 3260
- SIGEV\_SIGNAL\_CODE 1073, 1078,  
1633, 2778, 3260
- SIGEV\_SIGNAL\_CODE\_INIT()* 2781
- SIGEV\_SIGNAL\_INIT()* 2780
- SIGEV\_SIGNAL\_THREAD 1073,  
1078, 1633, 2778, 3260
- SIGEV\_SIGNAL\_THREAD\_INIT()* 2781
- SIGEV\_SIGNAL\_VALUE\_INIT()* 2780
- SIGEV\_THREAD 2778
- SIGEV\_THREAD\_INIT()* 2782
- SIGEV\_UNBLOCK 1662, 2778
- SIGEV\_UNBLOCK\_INIT()* 2782
- sigfillset()* 2783
- SIGFPE 2797
- SIGHUP 563, 2796
  - process groups, targeting 564
- SIG\_IGN 517, 2766, 2791, 2799
- SIGILL 2796
- siginfo\_t** 2766, 2800, 2805
- SIGINT 2796
  - process groups, sending  
to 3197
- SIGIO 2797
  - default actions 2765, 2799
- SIGIOT 2796
- sigismember()* 2785
- SIGKILL 2206, 2767, 2772, 2792,  
2793, **2797**
- siglongjmp()* 2787
- sigmask()* 2789
- .\_SIGMAX 2798
- .\_SIGMIN 2798
- sign, copying 343
- signal()* 2791
- SignalAction()*, *SignalAction\_r()*  
2796
- SignalKill()*, *SignalKill\_r()* 2804
- SignalProcmask()*,  
*SignalProcmask\_r()* 2810
- signals
  - actions 2764, 2791, 2796
  - default 2765
  - blocking 2772, 2834
  - SIGCONT 2772
  - SIGKILL 2206, 2772
  - SIGSTOP 2206, 2772
  - SIGTTOU 685, 687, 2269,  
2271
  - catching
    - SIGKILL 2792
    - SIGSTOP 2792
  - ignoring 2766
    - SIGCHLD 884, 2793
    - SIGKILL 2767, 2793
    - SIGSTOP 2767, 2793
    - SIGTTOU 685, 687, 2269,  
2271
  - information about 2845
  - masks
    - constructing 2789

- restoring 2787
- saving 2832
- signal-blocked 2810
- threads 2206
- names 2796
- POSIX 2796
- process groups, targeting 564, 1416, 1419, 2767, 2801, 2804
- processes
  - pending 2824
  - queuing 2765, 2829
  - suspending until delivered 2813
- sending 1416, 1419, 2283, 2804
- SIGABRT 113
- SIGALRM 141, 826, 2689
- SIGBUS 1586
- SIGCHLD 564, 2765, 2792
- SIGCONT 564
- SIGHUP 563
- SIGKILL 2792
- SIGPIPE 685, 687, 905, 1611, 1613, 2269, 2271, 2273
- SIGSEGV 1078, 1762, 1766, 2024, 2720
- SIGTRAP 363
- SIGXCPU 2720
- SIGXFSZ 2720
- sets
  - adding to 2770
  - initializing 2776, 2783
  - membership, checking for 2785
  - removing from 2774
- string describing 3066
- threads
  - mask 2826, 2836
  - threads, targeting 2116, 2767, 2801, 2804
  - unblocking 2841
  - user-defined 2797
  - waiting for 1878, 2816, 2822, 2836, 2838, 2843, 2845
- \_\_signalstub()* 2795
- SignalSuspend()*, *SignalSuspend\_r()* 2813
- SignalWaitinfo()*, *SignalWaitinfo\_r()* 2816
- significand()*, *significandf()* 2819
- sigpause()* 2822
- sigpending()* 2824
- SIGPIPE 685, 687, 905, 1611, 1613, 2269, 2271, 2273, **2797**
- SIGPOLL 2797
- sigprocmask()* 2826
- SIGPWR 2797
- sigqueue()* 2829
- SIGQUIT 2796
- SIGRTMAX 2798
- SIGRTMIN 2798
- SIGSEGV 1078, 1100, 1102, 1762, 1766, 2024, 2720, **2797**
- sigsetjmp()* 2832
- SIG\_SETMASK 2826
- sigsetmask()* 2834
- SIGSTOP 2206, 2767, 2772, 2792, 2793, **2797**
  - default actions 2765
- sigsuspend()* 2836
- SIGSYS 2797
- SIGTERM 2797

- sigtimedwait()* 2838
- SIGTRAP 363, 2796
- SIGTSTP 2797
- SIGTTIN 2797, 3417
- SIGTTOU 685, 687, 2269, 2271,  
2797, 3417
- SIG\_UNBLOCK 2826
- sigunblock()* 2841
- SIGURG 2797
  - default actions 2765, 2799
  - process groups, sending  
to 2909
- SIGUSR1 2797
- SIGUSR2 2797
- sigwait()* 2843
- sigwaitinfo()* 2845
- SIGWINCH 2797
  - default actions 2765, 2799
- SIGXCPU 2720
- SIGXFSZ 2720, 3334
- SI\_MAXAVAIL 2780
- SIMESGQ 2767, 2800, 2806
- SIMINAVAIL 2780
- Simple Network Management  
Protocol *See* SNMP
- sin()*, *sinf()* 2847
- sines 2847
  - hyperbolic 2849
  - inverse hyperbolic 156
- sinh()*, *sinhf()* 2849
- SIOCGIFCONF 1316
- SIOCGIFFLAGS 1316
- SI\_QUEUE 2767, 2800, 2806
- S\_IREAD 1115, 2995
- S\_IRGRP 1115, 2995, 3353
- S\_IROTH 1115, 2995, 3353
- S\_IRUSR 1115, 2995, 3353
- S\_IRWXG 1115, 2634, 2995, 3353
- S\_IRWXO 1115, 2634, 2995, 3353
- S\_IRWXU 1115, 2634, 2995, 3353
- S\_ISBLK()* 710, 2996
- S\_ISCHR()* 710, 2996
- S\_ISDIR()* 710, 2996
- S\_ISFIFO()* 710, 2996
- S\_ISGID 261, 262, 266, 578, 581,  
711, 1424, 2996, 3334
- S\_ISLNK()* 710, 2997
- S\_ISNAM()* 2997
- S\_ISREG()* 711, 2997
- S\_ISSOCK()* 2997
- S\_ISUID 261, 262, 266, 578, 581,  
711, 1424, 2995, 3334
- S\_ISVTX 261
- site-local addresses 1052
- SI\_TIMER 2767, 2800, 2806
- SI\_USER 2767, 2800, 2806
- S\_IWGRP 1115, 2995, 3353
- S\_IWOTH 1115, 2995, 3353
- S\_IWRITE 1115, 2995
- S\_IWUSR 1115, 2995, 3353
- S\_IXGRP 1115, 2995, 3353
- S\_IXOTH 1115, 2995, 3353
- S\_IXUSR 1115, 2995, 3353
- sleep()* 2851
- sleeping
  - for microseconds 3385
  - for milliseconds 368, 1796,  
1797
- sleepon locks
  - destroying 2855
  - initializing 2857
  - locking 2210, 2859
  - unblocking 2208, 2212, 2853,  
2861



- unlocking 2218, 2863
- waiting 2214, 2220, 2865
- \_sleepon\_broadcast()* 2853
- \_sleepon\_destroy()* 2855
- \_sleepon\_init()* 2857
- \_sleepon\_lock()* 2859
- \_sleepon\_signal()* 2861
- \_sleepon\_unlock()* 2863
- \_sleepon\_wait()* 2865
- slogb()* 2867
- \_SLOG\_CRITICAL* 2870
- \_SLOG\_DEBUG1* 2870
- \_SLOG\_DEBUG2* 2870
- \_SLOG\_ERROR* 2870
- slogf()* 2869
- slogi()* 2873
- \_SLOG\_INFO* 2870
- \_SLOG\_NOTICE* 2870
- \_SLOG\_SETCODE()* 2869, 3433
- \_SLOG\_SHUTDOWN* 2870
- \_SLOG\_WARNING* 2870
- \_smalloc()* 2875
- SNMP (Simple Network Management Protocol)
  - classification 106
  - daemon, configuration file for 2311
  - messages
    - creating 2887
    - freeing 2879
    - reading 2889
    - sending 2894
  - Protocol Data Unit (PDU)
    - creating 2887
    - freeing 2879
    - processing 2889
    - sending 2894
  - structure 2883
  - sessions
    - characteristics 2897
    - closing 2877
    - opening 2881
    - timeouts, handling 2901
    - transactions,
      - asynchronous 2892
  - snmp\_pdu* 2883
  - snmp\_session* 2881, 2897
  - snmp\_close()* 2877
  - SNMPCONFIGFILE** 2313
  - snmpd\_conf\_data* 2311
  - snmpd.conf* 2311
  - SNMP\_DEFAULT\_ADDRESS* 2898
  - SNMP\_DEFAULT\_COMMUNITY\_LEN* 2897
  - SNMP\_DEFAULT\_ENTERPRISE\_LENGTH* 2884
  - SNMP\_DEFAULT\_ERRINDEX* 2884
  - SNMP\_DEFAULT\_ERRSTAT* 2884
  - SNMP\_DEFAULT\_PEERNAME* 2898
  - SNMP\_DEFAULT\_REMPORT* 2898
  - SNMP\_DEFAULT\_REQID* 2884
  - SNMP\_DEFAULT\_RETRIES* 2897
  - SNMP\_DEFAULT\_TIME* 2885
  - SNMP\_DEFAULT\_TIMEOUT* 2897
  - SNMPERR\_BAD\_ADDRESS* 2881, 2895
  - SNMPERR\_BAD\_LOCPORT* 2882
  - SNMPERR\_BAD\_SESSION* 2877, 2895
  - SNMPERR\_GENERR* 2882, 2888, 2895
  - snmp\_errno* 2877, 2881, 2888, 2895
  - snmp\_free\_pdu()* 2879
  - snmp\_open()* 2881

- snmp\_pdu\_create()* 2887
- snmp\_read()* 2889
  - using with *select()* 2892
- snmp\_select\_info()* 2892
  - using with *select()* 2892
- snmp\_send()* 2894
- snmp\_timeout()* 2901
  - using with *select()* 2892
- SNMP\_VERSION\_1 2883, 2899
- SNMP\_VERSION\_2 2883, 2899
- snprintf()* 2903
- SO\_BINDTODEVICE 904
- SO\_BROADCAST 904
- sockaddr\_un** 3377
- socketmark()* 2906
- sockcred** 3379
- SOCKCREDSIZE()* 3379
- SOCK\_DGRAM 323, 908, 1051, 1313, 1341, 2908, 2909, 3348, 3377, 3571
- socket()* 983, 985, 1051, 1313, 1320, 1338, 2480, 2594, 2909, 3189, 3348, 3377
- socketpair()* 2912
- sockets
  - addresses
    - errors 748
    - freeing 693
    - getting 753
    - structure 127
  - connections
    - accepting on 117, 2281
    - initiating 323, 1798, 2298
    - listening for 1447, 2475
    - shutting down 2762
    - status 1801
    - creating 2909
      - a pair of 2912
      - connected 2912
  - datagrams 2908, 2909, 3348
  - debugging 904
  - file descriptors, testing for
    - association 1360
  - messages
    - peeking at 2343, 2346, 2350
    - receiving from 2344, 2347, 2351
    - sending to 2651, 2653, 2657
  - names
    - binding to 206, 2293
    - getting 899, 2468
  - options
    - getting 902
    - setting 2730
  - out-of-band (OOB) mark 2906
  - privileged IP port, binding to 209
  - privileged ports, getting 2490
  - raw 2908, 2909
  - stream 2908, 2909
  - types 2908
    - determining 904, 908
- SOCK\_RAW 983, 985, 1313, 1314, 1320, 1338, 2908, 2909
- SOCKS 3571
  - classification 106
  - commands, executing
    - remotely 2488
  - compiling for 3571
  - initializing 2915
  - library 3572
  - sockets
    - connections 2281, 2298, 2475

- names 2293, 2468
- socks3r.lib** 3571
- SOCKSinit()* 2915, 3571
- SOCK\_STREAM 117, 323, 1051, 1313, 1341, 1448, 2281, 2296, 2908, 2909, 3189, 3377
- SO\_DEBUG 904
- SO\_DONTROUTE 904
- SO\_ERROR 904
- SO\_KEEPAVIVE 905
- SO\_LINGER 901, 905, 2729
- SOL\_SOCKET 901, 2482
- SO\_OOBINLINE 905
- sopen()* 2918
- sopenfd()* 2923
- SO\_RCVBUF 906
- SO\_RCVLOWAT 906
- SO\_RCVTIMEO 901, 906, 2729
- SO\_REUSEADDR 906
- SO\_REUSEPORT 907
- sorting
  - directory entries 147, 2531
  - quick sort 2277
- SO\_SNDBUF 906
- SO\_SNDLOWAT 907
- SO\_SNDTIMEO 901, 907, 2729
- SO\_TIMESTAMP 908
- SO\_TYPE 908
- SO\_USELOOPBACK 908, 2482
- space, amount free for a
  - filesystem 715
- space, filesystem 3002
- space, testing a character for 1374, 1397
- spawn()* 2927
- spawn\** family of functions 49, 57, 1843
- SPAWN\_CHECK\_SCRIPT 2925, 2953
- spawnl()* 2934
- spawnle()* 2939
- spawnlp()* 2944
- spawnlpe()* 2948
- SPAWN\_NEWPGROUP 2926, 2953
- spawnp()* 2954
- SPAWN\_SEARCH\_PATH 2926, 2953
- SPAWN\_SETGROUP 2926, 2929, 2953, 2955
- SPAWN\_SETND 2926, 2953
- SPAWN\_SETSIGDEF 2926, 2927, 2930, 2953, 2954, 2956
- SPAWN\_SETSIGMASK 2926, 2929, 2953, 2956
- spawnv()* 2961
- spawnve()* 2966
- spawnvp()* 2970
- spawnvpe()* 2974
- SPD (Security Policy Database) 1324
- special characters 1612
- spinlocks
  - destroying 2224
  - initializing 2226
  - locking 2228, 2230
  - unlocking 2232
- sporadic scheduling 2031
- sprintf()* 2978
- spwd** 2260
- sqrt()*, *sqrtf()* 2980
- square roots 2980
- srand()* 2982
- srand48()* 2984

- srandom()* 2986
- \_srealloc()* 2988
- sscanf()* 2991
- SS\_REPL\_MAX 2555
- stack
  - maximum size 2720
  - memory, allocating from 144
  - overflow, protecting
    - against 2024
  - threads 2013, 2015, 2017, 2035, 2037, 2039
- stat** 710, 2993
- stat()*, *stat64()* 1485, 2993
  - resource managers, implementing in 1271, 1273
- STATE\_CONDVAR 3109, 3299
- STATE\_INTR 3299
- STATE\_JOIN 3299
- STATE\_MUTEX 3108, 3109, 3299
- STATE\_RECEIVE 3299
- STATE\_REPLY 1711, 3299
- STATE\_SEM 3299
- STATE\_SEND 1711, 3299
- STATE\_SIGSUSPEND 3299
- STATE\_SIGWAITINFO 3299
- st\_atime* 1563, 1566, 1947, 2308, 2336
- statvfs**, **statvfs64** 714, 3001
- statvfs()*, *statvfs64()* 3001
- st\_ctime* 266, 579, 581, 1424, 1438, 1563, 1566, 1947, 2477, 3552
- stderr* 69, 586, 626, 1945, 3005
  - buffering 2662, 2696
  - command-line options, errors when parsing 843
- daemons 1978
- formatted messages on 504, 3414, 3450, 3472
- host errors 958
- STDERR\_FILENO 626, 3005
- stdin* 69, 586, 626, 3006
  - characters, reading 611, 764, 766
  - daemons 1978
  - input, formatted 2533, 3430, 3454, 3561
  - strings, reading 889
  - wide characters, reading 938
- STDIN\_FILENO 626, 1947, 1956, 3006
- stdout* 69, 586, 626, 2258, 3007
  - buffering 2662, 2696
  - characters, writing 680, 2251, 2253
  - daemons 1978
  - output, formatted 1968, 3428, 3452, 3548
  - strings, writing 2258
  - wide characters, writing 2270
- STDOUT\_FILENO 626, 1947, 1956, 3007
- st\_mtime* 1566, 1947
- st\_mode* 2996
- st\_mtime* 1438, 1563, 1566, 1947, 2477, 3552
- ST\_NOSUID 518, 715, 2930, 3002
- straddstr()* 3008
- strcasemp()* 3010
- strcat()* 3013
- strchr()* 3015
- strcmp()* 3017
- strcmpi()* 3019

- strcoll()* 3021, 3090
- strcpy()* 3023
- strcspn()* 3025
- ST\_RDONLY 715, 3002
- strdup()* 3027
- stream I/O
  - buffering
    - associating with 2660, 2742
    - block 2662
    - line 2696
  - characters
    - pushing back 3373, 3375
    - reading 609, 611, 615, 760, 762, 764, 766, 936, 938
    - writing 678, 680, 2247, 2249, 2251, 2253, 2268, 2270
  - closing 584, 586
  - end-of-file 602
    - clearing 277
  - errors 604
    - clearing 277
    - messages, printing 1945
  - file descriptors
    - associating streams with 599
    - getting 626
  - files
    - flushing 606, 641
    - locking 637, 729
    - opening 650
    - reading 1066
    - unlocking 734
  - input, formatted 703, 746, 2533, 3421, 3426, 3430, 3454, 3561
  - output, formatted 676, 741, 1968, 3418, 3424, 3428, 3452, 3548
  - pipes
    - closing 1943
    - creating 1947
    - opening 1955
  - position
    - getting 613, 720
    - setting 706, 708
  - reading 688
  - reopening 697
  - rewinding 2462
  - seeking 706
  - strings
    - reading 889
    - writing 682, 686, 2258
  - telling 720
  - temporary files 1573, 1575, 3315
  - wide characters
    - orientation 739
    - reading 621, 623
    - writing 684
  - words
    - getting next 934
    - writing 2266
  - writing 743
- stream sockets 2908, 2909
- stream, returning to remote
  - command 2295, 2488
- streams
  - flushing 3173
- strerror()* 3029
- strftime()* 3031
- stricmp()* 3037

- strings *See also* characters; wide characters
  - character, filling with 3054, 3064
  - comparing 202, 1536, 1542
    - case-insensitive 3010, 3019, 3037, 3043, 3052
    - case-sensitive 3017, 3048
    - locale's collating sequence, using 3021
  - concatenating 3008, 3013, 3046
  - configuration, getting and setting 318
  - copying 204, 1532, 3023, 3027, 3050, 3090
  - encrypting 353, 460, 2275, 2694
  - error messages 1945, 3029
  - formatted 2903, 2978, 3435, 3438
  - hexadecimal numbers, converting to/from 172
  - input, formatted 703, 2533, 3421, 3430
  - IP addresses, converting to/from 1036, 1039
  - IPv4 addresses, converting to/from 1015, 1017, 1032, 1034
  - length 3039
  - lowercase, converting to 3041
  - matching 647
  - network numbers, converting to/from 1030
  - node descriptors, converting to/from 1806, 1814
  - numbers, converting to/from 170, 172, 174, 176, 1405, 1488, 3072, 3076, 3083, 3086, 3350, 3397
  - output, formatted 676, 1968, 3418, 3428
  - paths, resolving 2341
  - reversing 3060
  - scanning input from 2991, 3441, 3446
  - searching
    - characters 1013, 2470, 3015, 3058
    - sets of characters 3025, 3056, 3068
    - sets of wide characters 3486, 3498, 3504
    - slashes (/) 199, 385
    - strings 3070
    - wide characters 3478, 3500
  - signal descriptions 3066
  - splitting 3062, 3077, 3080
  - stdin*, reading from 889
  - stdout*, writing to 2258
  - streams
    - reading from 615
    - writing 682, 686
  - substrings
    - comparing,
      - case-insensitive 3043, 3052
      - comparing,
        - case-sensitive 3048
  - time\_t**, converting to/from 357
  - time, formatted 3031, 3488

- tm, converting to/from 152
  - tokens, splitting into 3077, 3080
  - transforming 3090
  - uppercase, converting to 3088
  - zeroing 220
- strlen()* 3039
- strlwr()* 3041
- strncasecmp()* 3043
- strncat()* 3046
- strncmp()* 3048, 3090
- strncpy()* 3050, 3090
- strnicmp()* 3052
- strnset()* 3054
- strpbrk()* 3056
- strrchr()* 3058
- strrev()* 3060
- strsep()* 3062
- strset()* 3064
- strsignal()* 3066
- strspn()* 3068
- strstr()* 3070
- strtod()* 3072
- strtoimax()* 3076
- strtok()* 3077
- strtok\_r()* 3080
- strtol()*, *strtoll()* 3083
- strtoul()*, *strtoull()* 3086
- strtoumax()* 3076
- struct**, offset of members
  - within 1833
- strupr()* 3088
- strxfrm()* 3090
- S\_TYPEISMQ()* 711, 2997
- S\_TYPEISSEM()* 711, 2997
- S\_TYPEISSHM()* 711, 2997
- suboptions, parsing 918
- SUN\_LEN()* 3377
- swab()* 3093
- swprintf()* 3095
- swscanf()* 3097
- symbolic links
  - creating 3099
  - deleting 2368, 3380
  - information, getting 1485
  - ownership, changing 1423
  - reading 2331
  - resolving 2341
  - temporary 1874, 1876
- symlink()* 3099
- SYMLOOP\_MAX 2332
- sync()* 3102
  - resource managers,
    - implementing in 1278, 1280
- SyncCondvarSignal()*, *SyncCondvarSignal\_r()* 3104
- SyncCondvarWait()*, *SyncCondvarWait\_r()* 3107
- SyncCtl()*, *SyncCtl\_r()* 3112
- SyncDestroy()*, *SyncDestroy\_r()* 3114
- synchronization objects *See also*
  - mutexes; semaphores;
  - threads
    - creating 3133
    - destroying 3114
  - mutexes
    - events 3112, 3117
    - locking 3119, 3124
    - priority 3112
    - reviving 3122

- semaphores
  - incrementing 3127, 3129
- threads
  - blocking 3107
  - waking up 3104
- SyncMutexEvent()*,
  - SyncMutexEvent\_r()* 3117
- SyncMutexLock()*,
  - SyncMutexLock\_r()* 3119
- SyncMutexRevive()*,
  - SyncMutexRevive\_r()*  
3122
- SyncMutexUnlock()*,
  - SyncMutexUnlock\_r()*  
3124
- SyncSemPost()*, *SyncSemPost\_r()*  
3127
- SyncSemWait()*, *SyncSemWait\_r()*  
3129
- SyncTypeCreate()*,
  - SyncTypeCreate\_r()* 3133
- sysconf()* 3136
- sysctl()* 3140
- sys\_errlist* 507
- syslog()* 3147
- sysmgr\_reboot()* 3150
- sys\_nerr* 508
- sys\_nsig* 2796
- SYSPAGE\_CPU\_ENTRY()* 3152
- SYSPAGE\_ENTRY()* 3154
  - qtime* 3154
    - boot\_time* 3154
    - cycles\_per\_sec* 3154
  - \_syspage\_ptr* 3157
- sys\_siglist* 2796
- system
  - clock
    - getting 306
    - period, getting and setting 304
    - setting 306
    - ticks per second 3154
  - commands, executing 3158
  - daemons 360, 1978
    - termination, notification of 1980
  - events
    - notification of 1981
    - triggering 1985
  - hardware information 973,  
975, 977, 979
  - information, getting and setting 3140
  - instruction set architecture 318
  - limits, getting 3136
  - rebooting 3150
  - resources
    - creating 2501
    - destroying 2505
    - limits, getting 773, 881
    - limits, setting 2719
    - querying 2516
    - reserving 2494
    - returning 2507
    - usage, getting 884
  - time since booting 3154
  - time, adjusting 282
- System Analysis Toolkit (SAT) 1093, 3331
- system databases
  - groups
    - closing 463
    - ID, getting information about 786, 788



- membership 1061
- name, getting information
  - about 791, 793
- next entry, getting 783
- rewinding 2678
- passwords
  - closing 491
  - encrypting 353, 2275
  - entry, getting for a user 871, 873, 876, 878
  - entry, getting next 868
  - rewinding 2713
- shadow passwords
  - closing 493
  - entry, reading 618, 911, 915
  - entry, structure 2260
  - entry, writing 2260
  - rewinding 2732
- system message log
  - closing 314
  - log priority mask 2701
  - opening 1850
  - writing to 3147, 3448
    - blocks 2867
    - formatted output 2869, 3433
    - integers 2873
- system packet forwarding database
  - See* ROUTE
- system page 3157
  - CPU-specific entry, getting a pointer to 3152
  - entry, getting a pointer to 3154
- system()* 49, 3158

## T

- tan()*, *tanf()* 3161
- tangents 3161
  - hyperbolic 3163
  - inverse hyperbolic 165
- tanh()*, *tanhf()* 3163
- target operating system 111
- tcdrain()* 3165
- tcdropline()* 3167
- tcflow()* 3170
- tcflush()* 3173
- tcgetattr()* 3176
- tcgetpgrp()* 3178
- tcgetsid()* 3180
- tcgetsize()* 3182
- TCIFLUSH 3173
- tcinject()* 3184
- TCIOFF 3170
- TCIOFFHW 3171
- TCIOFLUSH 3173
- TCION 3171
- TCIONHW 3171
- tcischars()* 3187
- TCOFLUSH 3173
- TCOOFF 3170
- TCOOFFHW 3170
- TCOON 3170
- TCOONHW 3170
- TCP (Transmission Control Protocol) 3189
  - connection, closing 464
  - SOCKS 3572
- TCP/IP
  - address information
    - addrinfo** 127
  - addresses

- manipulating 1021
- network numbers,
  - extracting 1026
- strings, converting
  - to/from 1015, 1017, 1032, 1034, 1036, 1039
- errors 953, 958, 967
- host entries
  - getting 806, 810, 812, 815
- hosts database
  - opening 2682
- Internet domain names
  - compressing 427
  - expanding 429
- messages
  - receiving 2344, 2347, 2351
  - sending 2651, 2653, 2657
- network database
  - closing 489
  - opening 2703
- protocols database
  - closing 490
  - opening 2711
- services database
  - closing 492
  - entry structure 2659
  - opening 2725
- sockets
  - ports, binding to 209
- TCP\_KEEPALIVE 3190
- TCP\_MAXSEG 3190
- TCP\_NODELAY 908, 3190
- TCSADRAIN 3194
- TCSAFLUSH 3194
- TCSANOW 3194
- tcsendbreak()* 3192
- tcsetattr()* 3194
- tcsetpgrp()* 3197
- tcsetsid()* 3200
- tcsetsize()* 3202
- tell()* 1483
- tell(), tell64()* 3204
- telldir()* 3207
- tempnam()* 3209
- temporary files
  - creating 3315
  - creating and opening 1573
  - name, generating 1575, 3209, 3318
- TERM** 274
- terminal control
  - characters, injecting 3184
  - communications line
    - break condition, asserting 3192
    - disconnecting 3167
  - draining 3165
  - flow control 3170
  - flushing 3173
  - process group ID 3180
    - getting 3178, 3197
  - size 3182, 3202
- terminals
  - attributes, setting 230
  - canonical input buffer 673, 1866
  - control
    - attributes 3176, 3194
    - structure 3211
  - controlling
    - making 3200
    - path name 355
  - file descriptor, testing for
    - association with 1354

- input speed 232, 242
- operating attributes 1123
- output speed 234, 245
- raw input buffer 673, 1866
- reading 2318
- TERMINFO** 274
- termios** 230, 2318, 3194, 3211
- text segment
  - beginning of 217
  - end of 515
- ThreadCancel()*, *ThreadCancel\_r()* 3234
- ThreadCreate()*, *ThreadCreate\_r()* 3239
- ThreadCtl()*, *ThreadCtl\_r()* 1069, 1078, 1083, 1085, 1087, 1095, 1097, 1100, 1102, 3245
- ThreadDestroy()*, *ThreadDestroy\_r()* 3249
- ThreadDetach()*, *ThreadDetach\_r()* 3252
- ThreadJoin()*, *ThreadJoin\_r()* 3254
- thread\_pool\_attr\_t* 3219
- thread\_pool\_control()* 3216
- THREAD\_POOL\_CONTROL\_HIWATER 3215
- THREAD\_POOL\_CONTROL\_INCREMENT 3215
- THREAD\_POOL\_CONTROL\_LOWATER 3215
- THREAD\_POOL\_CONTROL\_MAXIMUM 3216
- THREAD\_POOL\_CONTROL\_NONBLOCK 3216
- thread\_pool\_create()* 3218
- thread\_pool\_destroy()* 3225
- thread\_pool\_limits()* 3229
- thread\_pool\_start()* 3231
- threads
  - aborting 1995
  - attributes 3239
    - contention scope 2011, 2033
    - destroying 1999
    - detach state 2001, 2022
    - guard area, size of 2003, 2024
    - initializing 2019
    - scheduling parameters 2007, 2029
    - scheduling policy 2009, 2031
    - stack address 2013, 2035
    - stack size 2017, 2039
    - stack, lazy 2015, 2037
  - barriers
    - attributes 2043, 2047, 2051
    - attributes, process-shared 2049, 2053
    - destroying 2041
    - initializing 2043
    - waiting at 2045
  - blocking 3107
  - busy-waiting 1784, 1789, 1791, 1793
  - calibrating 1786
  - canceling 2055, 3234
  - cancellation
    - cleanup handlers 2057, 2059
    - points 3241
    - points, creating 2234
    - state 2196
    - type 2198

- clock ID 2102
- concurrency 2100, 2200
- condition variables
  - attributes 2077, 2083
  - attributes, clock 2079, 2085
  - attributes,
    - process-shared 2081, 2087
  - blocking on 2070, 2074
  - destroying 2064, 3114
  - initializing 2066, 3133
  - unblocking 2062, 2068
- creating 2090, 3239
- data 2106, 2110, 2114, 2204, 3242
- destroying 3249
- detached 2022, 3241
- detaching from a
  - process 2094, 3252
- errno* 507, 3242
- files, locking 637
- fork handlers, registering 1997
- freezing 3247
- I/O priority, requesting 3246
- IDs
  - calling thread 2195
  - comparing 2096
- initializing 2163
- joinable 2022, 3241
- joining 2108, 3254
  - with a time limit 2235
- keys 2110, 2114
- local storage 2106, 2110, 2114, 2204, 3242
- misaligned access
  - response 3245
- mutexes
  - attributes, destroying 2137
  - attributes, initializing 2150
  - attributes, priority
    - ceiling 2139, 2152
  - attributes,
    - process-shared 2143, 2156
  - attributes, recursive 2145, 2158
  - attributes, scheduling
    - protocol 2141, 2154
  - attributes, type 2147, 2161
  - destroying 2118, 3114
  - initializing 2122, 3133
  - locking 2124, 2130, 2133
  - priority ceiling 2120, 2128
  - unlocking 2135
- once-initialization 2163
- pool *See also* resource managers
  - attributes, changing 3216, 3229
  - creating 3218
  - destroying 3225
  - starting 3231
- private data 2106, 2110, 2114, 2204, 3242
- processor affinity 3246
- read-write locks
  - attributes, creating 2191
  - attributes, destroying 2187
  - attributes,
    - process-shared 2189, 2193
  - destroying 2166, 3114
  - initializing 2168, 3133

- locking for reading 2171, 2173, 2179
- locking for writing 2176, 2181, 2185
- unlocking 2183
- return status 3240
- scheduling parameters 2104, 2203, 2570, 2577
- scheduling policy 2104, 2203, 2570, 2577, 3240
  - inheriting 2005, 2027, 3241
- scope 3241
- signal mask
  - getting 2826
  - restoring 2787
  - saving 2832
  - setting 2826, 2836
  - signal-blocked 2810
- signals
  - initial state 3242
  - mask 2206
  - targeting 2767, 2801, 2804
  - terminating on 3241
- signals, sending 2116
- sleepon locks
  - destroying 2855
  - initializing 2857
  - locking 2210, 2859
  - unblocking 2208, 2212, 2853, 2861
  - unlocking 2218, 2863
  - waiting 2214, 2220, 2865
- spinlocks
  - destroying 2224
  - initializing 2226
  - locking 2228, 2230
  - unlocking 2232
- stack 3239
- suspending 294, 1782, 2836, 2851
- synchronizing 2045
- terminating 2098
- terminating
  - unconditionally 1995
- unfreezing 3247
- waking up 2208, 2212, 2853, 2861, 3104
- yielding 2567, 2579
- zombies 3241
- ticksize, getting and setting 304
- `time_t` 3257
  - `tm`, converting to/from 949, 951
- time members
  - in attribute structure of resource managers 1136
- Time Stamp Counter (TSC) 284
- time to live (TTL) 1313
- time zone
  - abbreviations 3341
  - default 319
  - offset from UTC 3312
  - setting 3342
- `time()` 3257
- `timeb` 723
- TIMED\_OUT 2899, 2901
- timeout, setting on a blocking state 3275
- timeouts, SNMP 2901
- TIMER\_ABSTIME 293, 3271
- `TimerAlarm()`, `TimerAlarm_r()` 3282
- `timer_create()` 3259

- TimerCreate()*, *TimerCreate\_r()*
  - 3285
- timer\_delete()* 3263
- TimerDestroy()*, *TimerDestroy\_r()*
  - 3288
- timer\_getexpstatus()* 3265
- timer\_getoverrun()* 3267
- timer\_gettime()* 3269
- TimerInfo()*, *TimerInfo\_r()* 3291
- timers
  - alarm, scheduling 3282
  - creating 3285
  - destroying 3288
  - information about,
    - getting 3291
  - interval
    - value, getting 825
    - value, setting 2688
  - realtime
    - creating 3259
    - destroying 3263
    - expiry status 3265
    - overruns 3267
    - time until expiry 3269,
      - 3272, 3295
  - threads 294
  - timeout, setting on kernel
    - blocking state 3299
- timer\_settime()* 3272
- TimerSettime()*, *TimerSettime\_r()*
  - 3295
- TimerTimeout()*, *TimerTimeout\_r()*
  - 3299
- timer\_timeout()*, *timer\_timeout\_r()*
  - 3275
- times
  - booting, since 3154
- calendar
  - current 3257
  - local, converting
    - to/from 1454, 1456, 1577
  - structure 3313
- clock
  - adjusting 282
  - cycles 284
  - getting 290
  - getting and setting 306
  - ID, getting 286, 300
  - period, getting and
    - setting 304
  - resolution, getting 288
  - setting 297
- current
  - calendar 3257
  - getting 723, 923
  - setting 2735
- daylight saving time 362, 3342
- difference, calculating 380
- files
  - access 1283, 1295, 1298,
    - 3390
  - modification 736, 1283,
    - 1295, 1298, 3387, 3390
  - status-change 1283, 1295,
    - 1298
- formatting 2699, 3031, 3488
- local
  - calendar, converting
    - to/from 1454, 1456, 1577
- nanoseconds
  - timespec**, converting
    - to/from 1827, 3310
- processes

- execution time limit,
      - getting 2559
    - execution time, in clock
      - ticks 279
    - specification structure 3309
  - time\_t**
    - strings, converting
      - to/from 357
    - tm**, converting to/from 949, 951
  - timespec**
    - nanoseconds, converting
      - to/from 1827, 3310
  - tm**
    - strings, converting
      - to/from 152
  - times()* 3306
  - timespec** 288, 3309
    - nanoseconds, converting
      - to/from 1827, 3310
  - timespec2nsec()* 3310
  - timezone* 3312, 3342
  - \_TLS* 3242
  - tm** 949, 1454, 1577, 3313, 3488
    - strings, converting
      - to/from 152, 357
    - time\_t**, converting to/from
      - 949, 951
  - tmpfile()* 3315
  - TMP\_MAX** 3209
  - tmpnam()* 3318
  - tms** 3306
  - tokens, breaking a string into 3077, 3080, 3513
  - tolower()* 3321
  - TOS (type of service) 1313
  - TOSTOP** 3213
  - toupper()* 3323
  - towctrans()* 3325
  - tolower()* 3327
  - toupper()* 3329
  - T\_PTR** 2377, 2380, 2383, 2386
  - TraceEvent()* 3331
  - Transmission Control Protocol *See* TCP
  - trigonometry *See also* hyperbolic functions
    - arccosine 123
    - arcsine 154
    - arctangent 161, 163
    - cosine 345
    - sine 2847
    - tangent 3161
  - TRP2\_REQ\_MSG** 2887
  - TRP\_REQ\_MSG** 2887
  - truncate()* 3334
  - TRY\_AGAIN** 954, 959
  - TSC (Time Stamp Counter) 284
  - TTL (time to live) 1313
  - ttynam()* 3337
  - ttynam\_r()* 3339
  - type of service (TOS) 1313
  - TZ** 274, 3342
  - tzname* 3341, 3342
  - tzset()* 1577, 3034, 3342
- U**
- ualarm()* 3345
  - UDP (User Datagram Protocol) 3348
    - not supported by SOCKS 3571

- UIO\_MAXIOV 2334, 2336, 3558
  - ulltoa()* 3350
  - ultoa()* 3350
  - umask()* 3353
  - umount()* 3356
  - UNALIGNED\_PUT16()* 3358
  - UNALIGNED\_PUT32()* 3360
  - UNALIGNED\_PUT64()* 3362
  - UNALIGNED\_RET16()* 3364
  - UNALIGNED\_RET32()* 3366
  - UNALIGNED\_RET64()* 3368
  - uname()* 3370
  - ungetc()* 3373
  - ungetwc()* 3375
  - Unicode 97
  - union**, offset of members
    - within 1833
  - Unix classification 106
  - UNIX-domain protocol 3377
  - unlink()* 3380
  - unsetenv()* 3383
  - uppercase
    - characters, converting to 3323
    - strings, converting to 3088
    - testing a character for 1377, 1399
    - wide characters, converting to 3325, 3329, 3531
  - usage of system resources 884
  - User Datagram Protocol (UDP) 3348
    - not supported by SOCKS 3571
  - user information file
    - closing 494
    - entry 3393
    - reading 927, 932
    - renaming 3395
    - returning to beginning of 2740
    - searching 929
    - writing 2263
  - USER\_PROCESS 929, 3394
  - users
    - IDs
      - effective 779, 2672, 2717, 2737
      - real 925, 2717, 2737
      - saved 2737
      - set-user 3137
    - names 827, 829
    - password database, getting entry for 871, 873, 876, 878
    - processes, maximum per real user ID 3136
  - usleep()* 3385
  - utilities, locating 319
  - utimbuf** 736, 3387
  - utime()* 3387
    - resource managers, implementing in 1295, 1298
  - utimes()* 3390
  - utmp** 3393
  - utmpname()* 3395
  - utoa()* 3397
  - utsname** 3370
- ## V
- va\_arg()* 3400
  - va\_copy()* 3406
  - va\_end()* 3408
  - valloc()* 3412



- variable\_list** 2885
  - variable-length argument lists
    - ("varargs") 3400, 3406, 3408, 3410
    - coercion 3401
  - variables, global
    - \_amblksiz* 149
    - \_argc* 150
    - \_argv* 151
    - \_auxv* 198
    - \_btext* 217
    - daylight* 362
    - \_edata* 459
    - \_end* 462
    - errno* 507
    - \_etext* 515
    - optarg* 842
    - opterr* 843
    - optind* 842
    - optopt* 843
    - \_progname* 1993
    - stderr* 3005
    - stdin* 3006
    - stdout* 3007
    - sys\_errlist* 507
    - sys\_nerr* 508
    - sys\_nsig* 2796
    - \_syspage\_ptr* 3157
    - sys\_siglist* 2796
    - timezone* 3312
    - tzname* 3341
  - va\_start()* 3410
  - verr()*, *verrx()* 3414
  - vfork()* 3416
  - vfprintf()* 3418
  - vfscanf()* 3421
  - vwprintf()* 3424
  - vfwscanf()* 3426
  - video memory, sharing 1587
  - virtual 8086 mode 1107
  - void** pointers, size of 111
  - vprintf()* 3428
  - vscanf()* 3430
  - vslogf()* 3433
  - vsnprintf()* 3435
  - vsprintf()* 3438
  - vsscanf()* 3441
  - vswprintf()* 3444
  - vswscanf()* 3446
  - vsyslog()* 3448
  - vwarn()*, *vwarnx()* 3450
  - vwprintf()* 3452
  - vwscanf()* 3454
- ## W
- wait()* 3307, 3456, 3463, 3470
  - wait3()* 3460
  - wait4()* 3463
  - waitid()* 3467
  - waitpid()* 3307, 3470
  - warn()*, *warnx()* 3472
  - warnings, formatted on *stderr*
    - 3450, 3472
  - WCONTINUED 3459, 3462, 3466, 3469
  - WCOREDUMP() 3457
  - wcrtomb()* 3474
  - wcscat()* 3476
  - wcschr()* 3478
  - wcscmp()* 3480
  - wcscoll()* 3482, 3524

- wcscpy()* 3484
- wcscspn()* 3486
- wcsftime()* 3488
- wcslen()* 3490
- wcsncat()* 3492
- wcsncmp()* 3494
- wcsncpy()* 3496
- wcspbrk()* 3498
- wcsrchr()* 3500
- wcsrtombs()* 3502
- wcsspn()* 3504
- wcsstr()* 3506
- wcstod()* 3508
- wcstof()* 3508
- wcstoimax()* 3512
- wcstok()* 3513
- wcstol()* 3516
- wcstold()* 3508
- wcstoll()* 3516
- wcstombs()* 3518
- wcstoul(), wcstoull()* 3522
- wcstoumax()* 3512
- wctob()* 3526
- wctomb()* 3528
- wctrans()* 3531
- wctype()* 3533
- WEXITED 3459, 3462, 3466, 3469
- WEXITSTATUS() 3159, 3456
- whitespace, testing a character
  - for 1374, 1397
- wide characters *See also*
  - characters; strings
  - classes 3533
  - converting 3325, 3531
  - copying 3539
  - overlapping objects 3541
  - lowercase, converting to 3325, 3327, 3531
  - multibyte characters, converting to/from 1511, 1521, 3474, 3528
  - conversion object, status of 1514
  - searching
    - in a string 3478, 3500
    - in memory 3535
  - sets of, searching for 3486, 3498, 3504
  - setting 3543
  - single-byte characters,
    - converting to/from 218, 3526
  - stdin*, reading from 938
  - stdout*, writing to 2270
  - streams
    - orientation 739
    - pushing back 3375
    - reading 621, 623, 936
    - writing to 684, 2268
  - strings
    - comparing 3480, 3482, 3494, 3537
    - concatenating 3476, 3492
    - copying 3484, 3496, 3524
    - formatted 3095, 3444
    - input, formatted 746, 3097, 3426, 3454, 3561
    - length 3490
    - multibyte characters,
      - converting to/from 1516, 1518, 3502, 3518

- numbers, converting
    - to/from 3508, 3512, 3516, 3522
  - output, formatted 741, 3424, 3452, 3548
  - searching for a set of wide characters 3486, 3498, 3504
  - searching for a string 3506
  - searching for a wide character 3478, 3500
  - splitting into tokens 3513
  - streams, writing to 686
  - transforming 3524
  - testing for
    - alphabetic 1381
    - alphanumeric 1379
    - character class 1385
    - control character 1383
    - decimal digit 1387
    - hexadecimal digit 1401
    - lowercase 1391
    - printable 1389, 1393
    - punctuation 1395
    - uppercase 1399
    - whitespace 1397
  - uppercase, converting to 3325, 3329, 3531
  - WIFCONTINUED()* 3457
  - WIFEXITED()* 3457
  - WIFSIGNALED()* 3457
  - WIFSTOPPED()* 3457
  - wmemchr()* 3535
  - wmemcmp()* 3537
  - wmemcpy()* 3539
  - wmemmove()* 3541
  - wmemset()* 3543
  - WNOHANG 3459, 3462, 3466, 3469
  - WNOWAIT 3459, 3462, 3466, 3469
  - W\_OK 120, 456
  - word expansions 3545, 3547
  - wordexp()* 3545
  - wordfree()* 3547
  - working directory 258, 768, 940
  - wprintf()* 3548
  - WRDE\_NOSYS 3545
  - write()* 1135, 3550
    - resource managers, implementing in 1301
  - writeblock()* 3555
  - writev()* 3558
  - wscanf()* 3561
  - WSTOPPED 3459, 3463, 3466, 3470
  - WSTOPSIG()* 3457
  - WTERMSIG()* 3457
  - WUNTRACED 3459, 3463, 3466, 3470
- X**
- X\_OK 120, 456
- Y**
- y0(), y0f()* 3563
  - y1(), y1f()* 3565
  - yn(), ynf()* 3567

## Z

zombies

- preventing children from
  - becoming 2766, 2799
- threads 3241