

AIX Version 6.1



# Technical Reference: Communications, Volume 2



AIX Version 6.1



# Technical Reference: Communications, Volume 2

**Note**

Before using this information and the product it supports, read the information in “Notices,” on page 453.

**Third Edition (October 2009)**

This edition applies to AIX Version 6.1 and to all subsequent releases of this product until otherwise indicated in new editions.

A reader's comment form is provided at the back of this publication. If the form has been removed, address comments to Information Development, Department 04XA-905-6B013, 11501 Burnet Road, Austin, Texas 78758-3400. To send comments electronically, use this commercial Internet address: [pserinfo@us.ibm.com](mailto:pserinfo@us.ibm.com). Any information that you supply may be used without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1997, 2009.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>About This Book</b>	xi
Highlighting	xi
Case-Sensitivity in AIX	xi
ISO 9000	xi
32-Bit and 64-Bit Support for the Single UNIX Specification	xii
Related Publications	xii
 <b>Chapter 1. Simple Network Management Protocol (SNMP)</b>	 1
getsmuxEntrybyname or getsmuxEntrybyidentity Subroutine	1
isodetailor Subroutine	2
ll_hdinit, ll_dbinit, ll_log, or ll_log Subroutine	3
o_number, o_integer, o_string, o_igeneric, o_generic, o_specific, or o_ipaddr Subroutine	5
oid_cmp, oid_cpy, oid_free, sprintoid, str2oid, ode2oid, oid2ode, oid2ode_aux, prim2oid, or oid2prim Subroutine	7
oid_extend or oid_normalize Subroutine	9
readobjects Subroutine	10
s_generic Subroutine	11
smux_close Subroutine	12
smux_error Subroutine	13
smux_free_tree Subroutine	14
smux_init Subroutine	15
smux_register Subroutine	16
smux_response Subroutine	17
smux_simple_open Subroutine	18
smux_trap Subroutine	20
smux_wait Subroutine	21
text2inst, name2inst, next2inst, or nexttot2inst Subroutine	22
text2oid or text2obj Subroutine	23
 <b>Chapter 2. Sockets</b>	 25
_getlong Subroutine	25
_getshort Subroutine	26
_putlong Subroutine	27
_putshort Subroutine	28
accept Subroutine	29
arpresolve_common Subroutine	30
arpupdate Subroutine	31
bind Subroutine	33
connect Subroutine	35
CreateloCompletionPort Subroutine	37
dn_comp Subroutine	38
dn_expand Subroutine	40
eaccept Subroutine	41
ebind Subroutine	42
econnect Subroutine	44
endhostent Subroutine	45
endhostent_r Subroutine	46
endnetent Subroutine	47
endnetent_r Subroutine	47
endnetgrent_r Subroutine	48
endprotoent Subroutine	49
endprotoent_r Subroutine	50
endservent Subroutine	50

endservent_r Subroutine . . . . .	51
erecv, erecvmsg, erecvfrom, enrecvmsg, or enrecvfrom Subroutine . . . . .	52
esend, esendto, or esendmsg Subroutine . . . . .	54
ether_ntoa, ether_aton, ether_ntohost, ether_hostton, or ether_line Subroutine . . . . .	56
FrcaCacheCreate Subroutine . . . . .	57
FrcaCacheDelete Subroutine . . . . .	58
FrcaCacheLoadFile Subroutine . . . . .	59
FrcaCacheUnloadFile Subroutine. . . . .	62
FrcaCtrlCreate Subroutine . . . . .	63
FrcaCtrlDelete Subroutine . . . . .	65
FrcaCtrlLog Subroutine . . . . .	66
FrcaCtrlStart Subroutine . . . . .	67
FrcaCtrlStop Subroutine . . . . .	68
freeaddrinfo Subroutine . . . . .	69
getaddrinfo Subroutine . . . . .	69
get_auth_method Subroutine . . . . .	72
getdomainname Subroutine. . . . .	73
gethostbyaddr Subroutine . . . . .	74
gethostbyaddr_r Subroutine. . . . .	75
gethostbyname Subroutine . . . . .	76
gethostbyname_r Subroutine . . . . .	78
gethostent Subroutine . . . . .	80
gethostent_r Subroutine . . . . .	80
gethostid Subroutine . . . . .	81
gethostname Subroutine . . . . .	82
GetMultipleCompletionStatus Subroutine . . . . .	83
getnameinfo Subroutine . . . . .	86
getnetbyaddr Subroutine . . . . .	87
getnetbyaddr_r Subroutine . . . . .	88
getnetbyname Subroutine . . . . .	89
getnetbyname_r Subroutine. . . . .	90
getnetent Subroutine . . . . .	91
getnetent_r Subroutine . . . . .	92
getnetgrent_r Subroutine. . . . .	93
getpeername Subroutine . . . . .	94
getprotobyname Subroutine. . . . .	95
getprotobyname_r Subroutine . . . . .	96
getprotobynumber Subroutine . . . . .	97
getprotobynumber_r Subroutine . . . . .	98
getprotoent Subroutine . . . . .	99
getprotoent_r Subroutine . . . . .	100
GetQueuedCompletionStatus Subroutine . . . . .	101
getservbyname Subroutine . . . . .	103
getservbyname_r Subroutine . . . . .	104
getservbyport Subroutine . . . . .	105
getservbyport_r Subroutine . . . . .	106
getservent Subroutine . . . . .	108
getservent_r Subroutine . . . . .	109
getsockname Subroutine . . . . .	109
getsockopt Subroutine . . . . .	111
htonl Subroutine . . . . .	119
htonll Subroutine . . . . .	119
htons Subroutine . . . . .	120
if_freenameindex Subroutine . . . . .	121
if_indextoname Subroutine . . . . .	122
if_nameindex Subroutine . . . . .	123

if_nametoindex Subroutine . . . . .	123
inet6_opt_append Subroutine . . . . .	124
inet6_opt_find Subroutine . . . . .	125
inet6_opt_finish Subroutine . . . . .	126
inet6_opt_get_val Subroutine. . . . .	126
inet6_opt_init Subroutine . . . . .	127
inet6_opt_next Subroutine . . . . .	128
inet6_opt_set_val Subroutine. . . . .	129
inet6_rth_add Subroutine . . . . .	129
inet6_rth_getaddr Subroutine. . . . .	130
inet6_rth_init Subroutine . . . . .	130
inet6_rth_reverse Subroutine. . . . .	131
inet6_rth_segments Subroutine . . . . .	132
inet6_rth_space Subroutine . . . . .	133
inet_addr Subroutine . . . . .	133
inet_lnaof Subroutine . . . . .	136
inet_makeaddr Subroutine. . . . .	137
inet_net_ntop Subroutine . . . . .	138
inet_net_pton Subroutine . . . . .	139
inet_netof Subroutine . . . . .	140
inet_network Subroutine . . . . .	141
inet_ntoa Subroutine . . . . .	143
inet_ntop Subroutine . . . . .	144
inet_pton Subroutine . . . . .	145
innetgr, getnetgrent, setnetgrent, or endnetgrent Subroutine . . . . .	146
ioctl Socket Control Operations . . . . .	147
isinet_addr Subroutine . . . . .	154
kvalid_user Subroutine . . . . .	156
listen Subroutine . . . . .	157
ntohl Subroutine . . . . .	158
ntohl Subroutine . . . . .	159
ntohs Subroutine . . . . .	160
PostQueuedCompletionStatus Subroutine . . . . .	161
rcmd Subroutine . . . . .	162
rcmd_af Subroutine . . . . .	164
ReadFile Subroutine . . . . .	165
recv Subroutine. . . . .	167
recvfrom Subroutine . . . . .	169
recvmsg Subroutine . . . . .	171
res_init Subroutine . . . . .	172
res_mkquery Subroutine . . . . .	173
res_ninit Subroutine . . . . .	176
res_query Subroutine . . . . .	178
res_search Subroutine . . . . .	180
res_send Subroutine . . . . .	182
rexec Subroutine . . . . .	183
rexec_af Subroutine . . . . .	185
rresvport Subroutine . . . . .	186
rresvport_af Subroutine . . . . .	187
ruserok Subroutine . . . . .	188
sctp_opt_info Subroutine . . . . .	190
sctp_peeloff Subroutine. . . . .	191
sctpctrl Subroutine . . . . .	192
send Subroutine . . . . .	193
sendmsg Subroutine . . . . .	195
sendto Subroutine. . . . .	197

send_file Subroutine . . . . .	199
set_auth_method Subroutine . . . . .	204
setdomainname Subroutine . . . . .	205
sethostent Subroutine . . . . .	206
sethostent_r Subroutine. . . . .	207
sethostid Subroutine . . . . .	208
sethostname Subroutine . . . . .	209
setnetent Subroutine . . . . .	210
setnetent_r Subroutine . . . . .	211
setnetgrent_r Subroutine . . . . .	211
setprotoent Subroutine . . . . .	212
setprotoent_r Subroutine . . . . .	213
setservent Subroutine . . . . .	214
setservent_r Subroutine. . . . .	215
setsockopt Subroutine . . . . .	216
shutdown Subroutine. . . . .	226
SLPAttrCallback Subroutine . . . . .	227
SLPClose Subroutine . . . . .	228
SLPDereg Subroutine . . . . .	228
SLPEscape Subroutine . . . . .	229
SLPFindAttrs Subroutine . . . . .	230
SLPFindScopes Subroutine . . . . .	231
SLPFindSrvs Subroutine . . . . .	232
SLPFindSrvTypes Subroutine . . . . .	233
SLPFree Subroutine . . . . .	234
SLPGetProperty Subroutine . . . . .	234
SLPOpen Subroutine . . . . .	235
SLPParseSrvURL Subroutine . . . . .	236
SLPReg Subroutine . . . . .	237
SLPRegReport Callback Subroutine . . . . .	238
SLPSrvTypeCallback Subroutine . . . . .	239
SLPSrvURLCallback Subroutine . . . . .	240
SLPUnescape Subroutine . . . . .	240
socket Subroutine . . . . .	241
socketpair Subroutine . . . . .	243
socks5_getserv Subroutine . . . . .	245
/etc/socks5c.conf File . . . . .	246
socks5tcp_accept Subroutine . . . . .	247
socks5tcp_bind Subroutine . . . . .	249
socks5tcp_connect Subroutine . . . . .	251
socks5udp_associate Subroutine . . . . .	253
socks5udp_sendto Subroutine . . . . .	256
splice Subroutine . . . . .	258
WriteFile Subroutine . . . . .	258
 <b>Chapter 3. Streams . . . . .</b>	 261
adjmsg Utility . . . . .	261
alloca Utility . . . . .	261
backq Utility . . . . .	262
bcanput Utility . . . . .	263
bufcall Utility . . . . .	263
canput Utility. . . . .	265
clone Device Driver . . . . .	265
copyb Utility . . . . .	266
copymsg Utility . . . . .	266
datamsg Utility . . . . .	267

dlpi STREAMS Driver . . . . .	268
dupb Utility . . . . .	269
dupmsg Utility . . . . .	269
enableok Utility . . . . .	270
esballoc Utility . . . . .	271
flushband Utility . . . . .	272
flushq Utility . . . . .	272
freeb Utility . . . . .	273
freemsg Utility . . . . .	274
getadmin Utility . . . . .	274
getmid Utility . . . . .	275
getmsg System Call . . . . .	275
getpmsg System Call . . . . .	278
getq Utility . . . . .	279
insq Utility . . . . .	280
ioctl Streams Device Driver Operations . . . . .	281
I_ATMARK streamio Operation . . . . .	292
I_CANPUT streamio Operation . . . . .	292
I_CKBAND streamio Operation . . . . .	293
I_FDINSERT streamio Operation . . . . .	293
I_FIND streamio Operation . . . . .	294
I_FLUSH streamio Operation . . . . .	295
I_FLUSHBAND streamio Operation . . . . .	295
I_GETBAND streamio Operation . . . . .	296
I_GETCLTIME streamio Operation . . . . .	296
I_GETSIG streamio Operation . . . . .	297
I_GRDOPT streamio Operation . . . . .	297
I_LINK streamio Operation . . . . .	297
I_LIST streamio Operation . . . . .	298
I_LOOK streamio Operation . . . . .	299
I_NREAD streamio Operation . . . . .	299
I_PEEK streamio Operation . . . . .	300
I_PLINK streamio Operation . . . . .	300
I_POP streamio Operation . . . . .	301
I_PUNLINK streamio Operation . . . . .	302
I_PUSH streamio Operation . . . . .	302
I_RECVFD streamio Operation . . . . .	303
I_SENDFD streamio Operation . . . . .	304
I_SETCLTIME streamio Operation . . . . .	304
I_SETSIG streamio Operation . . . . .	305
I_SRDOPT streamio Operation . . . . .	306
I_STR streamio Operation . . . . .	307
I_UNLINK streamio Operation . . . . .	308
isastream Function . . . . .	309
linkb Utility . . . . .	309
mi_bufcall Utility . . . . .	310
mi_close_comm Utility . . . . .	311
mi_next_ptr Utility . . . . .	312
mi_open_comm Utility . . . . .	313
msgdsize Utility . . . . .	314
noenable Utility . . . . .	315
OTHERQ Utility . . . . .	315
pfmod Packet Filter Module . . . . .	316
pullupmsg Utility . . . . .	319
putbq Utility . . . . .	320
putctl1 Utility . . . . .	320

putctl Utility . . . . .	321
putmsg System Call . . . . .	322
putnext Utility . . . . .	324
putpmsg System Call . . . . .	324
putq Utility . . . . .	325
qenable Utility . . . . .	327
qreply Utility . . . . .	327
qsize Utility . . . . .	328
RD Utility . . . . .	328
rmvb Utility . . . . .	329
rmvq Utility . . . . .	329
sad Device Driver . . . . .	330
splstr Utility . . . . .	333
splx Utility . . . . .	333
srv Utility . . . . .	333
str_install Utility . . . . .	335
streamio Operations . . . . .	339
strlog Utility . . . . .	341
strqget Utility . . . . .	342
t_accept Subroutine for Transport Layer Interface . . . . .	343
t_alloc Subroutine for Transport Layer Interface . . . . .	345
t_bind Subroutine for Transport Layer Interface . . . . .	347
t_close Subroutine for Transport Layer Interface . . . . .	349
t_connect Subroutine for Transport Layer Interface . . . . .	350
t_error Subroutine for Transport Layer Interface . . . . .	352
t_free Subroutine for Transport Layer Interface . . . . .	353
t_getinfo Subroutine for Transport Layer Interface . . . . .	354
t_getstate Subroutine for Transport Layer Interface . . . . .	357
t_listen Subroutine for Transport Layer Interface . . . . .	358
t_look Subroutine for Transport Layer Interface . . . . .	360
t_open Subroutine for Transport Layer Interface . . . . .	361
t_optmgmt Subroutine for Transport Layer Interface . . . . .	363
t_rcv Subroutine for Transport Layer Interface . . . . .	365
t_rcvconnect Subroutine for Transport Layer Interface . . . . .	366
t_rcvdis Subroutine for Transport Layer Interface . . . . .	368
t_rcvrel Subroutine for Transport Layer Interface . . . . .	370
t_rcvudata Subroutine for Transport Layer Interface . . . . .	371
t_rcvuderr Subroutine for Transport Layer Interface . . . . .	372
t_snd Subroutine for Transport Layer Interface . . . . .	373
t_snddis Subroutine for Transport Layer Interface . . . . .	375
t_sndrel Subroutine for Transport Layer Interface . . . . .	376
t_sndudata Subroutine for Transport Layer Interface . . . . .	377
t_sync Subroutine for Transport Layer Interface . . . . .	379
t_unbind Subroutine for Transport Layer Interface . . . . .	380
testb Utility . . . . .	381
timeout Utility . . . . .	382
timod Module . . . . .	383
tirdwr Module . . . . .	385
unbufcall Utility . . . . .	386
unlinkb Utility . . . . .	386
untimeout Utility . . . . .	387
unweldq Utility . . . . .	388
wantio Utility . . . . .	389
wantmsg Utility . . . . .	390
weldq Utility . . . . .	392
WR Utility . . . . .	393

xtiso STREAMS Driver . . . . .	393
t_accept Subroutine for X/Open Transport Interface . . . . .	395
t_alloc Subroutine for X/Open Transport Interface . . . . .	397
t_bind Subroutine for X/Open Transport Interface . . . . .	399
t_close Subroutine for X/Open Transport Interface . . . . .	402
t_connect Subroutine for X/Open Transport Interface . . . . .	403
t_error Subroutine for X/Open Transport Interface . . . . .	405
t_free Subroutine for X/Open Transport Interface . . . . .	406
t_getinfo Subroutine for X/Open Transport Interface . . . . .	408
t_getprotaddr Subroutine for X/Open Transport Interface . . . . .	410
t_getstate Subroutine for X/Open Transport Interface . . . . .	411
t_listen Subroutine for X/Open Transport Interface . . . . .	412
t_look Subroutine for X/Open Transport Interface . . . . .	414
t_open Subroutine for X/Open Transport Interface . . . . .	415
t_optmgmt Subroutine for X/Open Transport Interface. . . . .	419
t_rcv Subroutine for X/Open Transport Interface . . . . .	425
t_rcvconnect Subroutine for X/Open Transport Interface . . . . .	427
t_rcvdis Subroutine for X/Open Transport Interface. . . . .	429
t_rcvrel Subroutine for X/Open Transport Interface . . . . .	430
t_rcvudata Subroutine for X/Open Transport Interface. . . . .	431
t_rcvuderr Subroutine for X/Open Transport Interface . . . . .	433
t_snd Subroutine for X/Open Transport Interface . . . . .	434
t_snddis Subroutine for X/Open Transport Interface . . . . .	436
t_sndrel Subroutine for X/Open Transport Interface . . . . .	437
t_sndudata Subroutine for X/Open Transport Interface . . . . .	439
t_strerror Subroutine for X/Open Transport Interface . . . . .	440
t_sync Subroutine for X/Open Transport Interface . . . . .	441
t_unbind Subroutine for X/Open Transport Interface . . . . .	442
Options for the X/Open Transport Interface . . . . .	443
<b>Chapter 4. Packet Capture . . . . .</b>	<b>449</b>
Packet Capture Library Subroutines . . . . .	449
ioctl BPF Control Operations . . . . .	449
<b>Appendix. Notices . . . . .</b>	<b>453</b>
Trademarks . . . . .	454
<b>Index . . . . .</b>	<b>455</b>



---

## About This Book

This book provides experienced C programmers with complete detailed information about Simple Network Management Protocol, sockets, streams, and packet capture library subroutines for the AIX operating system. To use the book effectively, you should be familiar with commands, system calls, subroutines, file formats, and special files. This publication is also available on the documentation CD that is shipped with the operating system.

This book is part of the six-volume technical reference set, *AIX Version 6.1 Technical Reference*, that provides information on system calls, kernel extension calls, and subroutines in the following volumes:

- *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 1* and *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 2* provide information on system calls, subroutines, functions, macros, and statements associated with base operating system runtime services.
- *AIX Version 6.1 Technical Reference: Communications Volume 1* and *AIX Version 6.1 Technical Reference: Communications Volume 2* provide information on entry points, functions, system calls, subroutines, and operations related to communications services.
- *AIX Version 6.1 Technical Reference: Kernel and Subsystems Volume 1* and *AIX Version 6.1 Technical Reference: Kernel and Subsystems Volume 2* provide information about kernel services, device driver operations, file system operations, subroutines, the configuration subsystem, the communications subsystem, the low function terminal (LFT) subsystem, the logical volume subsystem, the M-audio capture and playback adapter subsystem, the printer subsystem, the SCSI subsystem, and the serial DASD subsystem.

---

## Highlighting

The following highlighting conventions are used in this book:

### **Bold**

Identifies commands, subroutines, keywords, files, structures, directories, and other items whose names are predefined by the system. Also identifies graphical objects such as buttons, labels, and icons that the user selects.

### *Italics*

Identifies parameters whose actual names or values are to be supplied by the user.

### Monospace

Identifies examples of specific data values, examples of text similar to what you might see displayed, examples of portions of program code similar to what you might write as a programmer, messages from the system, or information you should actually type.

---

## Case-Sensitivity in AIX

Everything in the AIX operating system is case-sensitive, which means that it distinguishes between uppercase and lowercase letters. For example, you can use the **ls** command to list files. If you type **LS**, the system responds that the command is "not found." Likewise, **FILEA**, **FiLea**, and **filea** are three distinct file names, even if they reside in the same directory. To avoid causing undesirable actions to be performed, always ensure that you use the correct case.

---

## ISO 9000

ISO 9000 registered quality systems were used in the development and manufacturing of this product.

---

## 32-Bit and 64-Bit Support for the Single UNIX Specification

Beginning with Version 5.2, the operating system is designed to support The Open Group's Single UNIX Specification Version 3 (UNIX 03) for portability of UNIX-based operating systems. Many new interfaces, and some current ones, have been added or enhanced to meet this specification, making Version 5.2 even more open and portable for applications, while remaining compatible with previous releases of AIX. To determine the proper way to develop a UNIX 03-portable application, you may need to refer to The Open Group's UNIX 03 specification, which can be accessed online or downloaded from <http://www.unix.org/> .

---

## Related Publications

The following books contain information about or related to application programming interfaces:

- *Operating system and device management*
- *Networks and communication management*
- *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*
- *AIX Version 6.1 Communications Programming Concepts*
- *AIX Version 6.1 Kernel Extensions and Device Support Programming Concepts*
- *AIX Version 6.1 Files Reference*

---

# Chapter 1. Simple Network Management Protocol (SNMP)

---

## getsmuxEntrybyname or getsmuxEntrybyidentity Subroutine

### Purpose

Retrieves SNMP multiplexing (SMUX) peer entries from the **/etc/snmpd.peers** file or the local **snmpd.peers** file.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
struct smuxEntry *getsmuxEntrybyname ( name)
char *name;
struct smuxEntry *getsmuxEntrybyidentity ( identity)
OID identity;
```

### Description

The **getsmuxEntrybyname** and **getsmuxEntrybyidentity** subroutines read the **snmpd.peers** file and retrieve information about the SMUX peer. The sample peers file is **/etc/snmpd.peers**. However, these subroutines can also retrieve the information from a copy of the file that is kept in the local directory. The **snmpd.peers** file contains entries for the SMUX peers defined for the network. Each SMUX peer entry should contain:

- The name of the SMUX peer.
- The SMUX peer object identifier.
- An optional password to be used on connection initiation. The default password is a null string.
- The optional priority to register the SMUX peer. The default priority is 0.

The **getsmuxEntrybyname** subroutine searches the file for the specified name. The **getsmuxEntrybyidentity** subroutine searches the file for the specified object identifier.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

### Parameters

<i>name</i>	Points to a character string that names the SMUX peer.
<i>identity</i>	Specifies the object identifier for a SMUX peer.

### Return Values

If either subroutine finds the specified SMUX entry, that subroutine returns a structure containing the entry. Otherwise, a null entry is returned.

### Files

<b>/etc/snmpd.peers</b>	Contains the SMUX peer definitions for the network.
-------------------------	---

## Related Information

List of Network Manager Programming References.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## isodetailor Subroutine

### Purpose

Initializes variables for various logging facilities.

### Library

ISODE Library (**libisode.a**)

### Syntax

```
#include <isode/tailor.h>
void isodetailor (myname, wantuser)
char * myname;
int  wantuser;
```

### Description

The ISODE library contains internal logging facilities. Some of the facilities need to have their variables initialized. The **isodetailor** subroutine sets default or user-defined values for the logging facility variables. The logging facility variables are listed in the **usr/lpp/snmpd/smux/isodetailor** file.

The **isodetailor** subroutine first reads the **/etc/isodetailor** file. If the *wantuser* parameter is set to 0, the **isodetailor** subroutine ignores the *myname* parameter and reads the **/etc/isodetailor** file. If the *wantuser* parameter is set to a value greater than 0, the **isodetailor** subroutine searches the current user's home directory (**\$HOME**) and reads a file based on the *myname* parameter. If the *myname* parameter is specified, the **isodetailor** subroutine reads a file with the name in the form **.myname\_tailor**. If the *myname* parameter is null, the **isodetailor** subroutine reads a file named **.isode\_tailor**. The **\_tailor** file contents must be in the following form:

```
#comment
<variable> : <value>  # comment
<variable> : <value>  # comment
<variable> : <value>  # comment
```

The comments are optional. The **isodetailor** subroutine reads the file and changes the values. The latest entry encountered is the final value. The subroutine reads **/etc/isodetailor** first and then the **\$HOME** directory, if told to do so. A complete list of the variables is in the **/usr/lpp/snmpd/smux/isodetailor** sample file.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

### Parameters

<i>myname</i>	Contains a character string describing the SNMP multiplexing (SMUX) peer.
<i>wantuser</i>	Indicates that the <b>isodetailor</b> subroutine should check the <b>\$HOME</b> directory for a <b>isodetailor</b> file if the value is greater than 0. If the value of the <i>wantuser</i> parameter is set to 0, the <b>\$HOME</b> directory is not checked, and the <i>myname</i> parameter is ignored.

## Files

`/etc/isodetailor`

Location of user's copy of the `/usr/lpp/snmpd/smux/isodetailor` file.

`/usr/lpp/snmpd/smux/isodetailor`

Contains a complete list of all the logging parameters.

## Related Information

The `ll_hdinit`, `ll_dbinit`, `_ll_log`, or `ll_log` subroutine.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## ll\_hdinit, ll\_dbinit, \_ll\_log, or ll\_log Subroutine

### Purpose

Reports errors to log files.

### Library

ISODE Library (`libisode.a`)

### Syntax

```
#include <isode/logger.h>
void ll_hdinit (lp, prefix)
register LLog * lp;
char * prefix;
void ll_dbinit (lp, prefix)
register LLog * lp;
char * prefix;
int _ll_log (lp, event, ap)
register LLog * lp;
int event;
va_list ap;
int ll_log ( va_alist)
va_dcl
```

### Description

The ISODE library provides logging subroutines to put information into log files. The **LLog** data structure contains the log file information needed to control the associated log. The SMUX peer provides the log file information to the subroutines.

The **LLog** structure contains the following fields:

```
typedef struct ll_struct
{
    char    *ll_file;      /* path name to logging file      */
    char    *ll_hdr;      /* text to put in opening line    */
    char    *ll_dhdr;     /* dynamic header - changes      */
    int     ll_events;     /* loggable events                */
    int     ll_syslog;     /* loggable events to send to syslog */
                                /* takes same values as ll_events */
    int     ll_msize;     /* max size for log, in Kbytes    */
                                /* If ll_msize < 0, then no checking */
    int     ll_stat;      /* assorted switches              */
    int     ll_fd;        /* file descriptor                */
} LLog;
```

The possible values for the `ll_events` and `ll_syslog` fields are:

<code>LLOG_NONE</code>	<code>0</code>	<code>/* No logging is performed</code>	<code>*/</code>
<code>LLOG_FATAL</code>	<code>0x01</code>	<code>/* fatal errors</code>	<code>*/</code>
<code>LLOG_EXCEPTIONS</code>	<code>0x02</code>	<code>/* exceptional events</code>	<code>*/</code>
<code>LLOG_NOTICE</code>	<code>0x04</code>	<code>/* informational notices</code>	<code>*/</code>
<code>LLOG_PDUS</code>	<code>0x08</code>	<code>/* PDU printing</code>	<code>*/</code>
<code>LLOG_TRACE</code>	<code>0x10</code>	<code>/* program tracing</code>	<code>*/</code>
<code>LLOG_DEBUG</code>	<code>0x20</code>	<code>/* full debugging</code>	<code>*/</code>
<code>LLOG_ALL</code>	<code>0xff</code>	<code>/* All of the above logging</code>	<code>*/</code>

The possible values for the `ll_stat` field are:

<code>LLOGNIL</code>	<code>0x00</code>	<code>/* No status information</code>	<code>*/</code>
<code>LLOGCLS</code>	<code>0x01</code>	<code>/* keep log closed, except writing</code>	<code>*/</code>
<code>LLOGCRT</code>	<code>0x02</code>	<code>/* create log if necessary</code>	<code>*/</code>
<code>LLOGZER</code>	<code>0x04</code>	<code>/* truncate log when limits reach</code>	<code>*/</code>
<code>LLOGERR</code>	<code>0x08</code>	<code>/* log closed due to (soft) error</code>	<code>*/</code>
<code>LLOGTTY</code>	<code>0x10</code>	<code>/* also log to stderr</code>	<code>*/</code>
<code>LLOGHDR</code>	<code>0x20</code>	<code>/* static header allocated/filled</code>	<code>*/</code>
<code>LLOGDHR</code>	<code>0x40</code>	<code>/* dynamic header allocated/filled</code>	<code>*/</code>

The **`ll_hdinit`** subroutine fills the `ll_hdr` field of the **`LLog`** record. The subroutine allocates the memory of the static header and creates a string with the information specified by the *prefix* parameter, the current user's name, and the process ID of the SMUX peer. It also sets the static header flag in the `ll_stat` field. If the *prefix* parameter value is null, the header flag is set to the "unknown" string.

The **`ll_dbinit`** subroutine fills the `ll_file` field of the **`LLog`** record. If the *prefix* parameter is null, the `ll_file` field is not changed. The **`ll_dbinit`** subroutine also calls the **`ll_hdinit`** subroutine with the same *lp* and *prefix* parameters. The **`ll_dbinit`** subroutine sets the log messages to **`stderr`** and starts the logging facility at its highest level.

The **`_ll_log`** and **`ll_log`** subroutines are used to print to the log file. When the **`LLog`** structure for the log file is set up, the **`_ll_log`** or **`ll_log`** subroutine prints the contents of the string format, with all variables filled in, to the log specified in the *lp* parameter. The **`LLog`** structure passes the name of the target log to the subroutine.

The expected parameter format for the **`_ll_log`** and **`ll_log`** subroutines is:

- **`_ll_log(lp, event, what, string_format, ...)`**;
- **`ll_log(lp, event, what, string_format, ...)`**;

The difference between the **`_ll_log`** and the **`ll_log`** subroutine is that the **`_ll_log`** uses an explicit listing of the **`LLog`** structure and the *event* parameter. The **`ll_log`** subroutine handles all the variables as a variable list.

The *event* parameter specifies the type of message being logged. This value is checked against the `events` field in the log record. If it is a valid event for the log, the other **`LLog`** structure variables are written to the log.

The *what* parameter variable is a string that explains what actions the subroutines have accomplished. The rest of the variables should be in the form of a **`printf`** statement, a string format and the variables to fill the various variable placeholders in the string format. The final output of the logging subroutine is in the following format:

```
mm/dd hh:mm:ss ll_hdr ll_dhdr string_format what: system_error
```

where:

Variable	Description
<code>mm/dd</code>	Specifies the date.

Variable	Description
hh:mm:ss	Specifies the time.
ll_hdr	Specifies the value of the ll_hdr field of the <b>LLog</b> structure.
ll_dhdr	Specifies the value of the ll_dhdr field of the <b>LLog</b> structure.
string_format	Specifies the string format passed to the <b>ll_log</b> subroutine, with the extra variables filled in.
what	Specifies the variable that tells what has occurred. The what variable often contains the reason for the failure. For example if the memory device, <b>/dev/mem</b> , fails, the what variable contains the name of the <b>/dev/mem</b> device.
system_error	Contains the string for the <b>errno</b> value, if it exists.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>lp</i>	Contains a pointer to a structure that describes a log file. The <i>lp</i> parameter is used to describe things entered into the log, the file name, and headers.
<i>prefix</i>	Contains a character string that is used to represent the name of the SMUX peer in the <b>ll_hdinit</b> subroutine. In the <b>ll_dbinit</b> subroutine, the <i>prefix</i> parameter represents the name of the log file to be used. The new log file name will be <i>.lprefix.log</i> .
<i>event</i>	Specifies the type of message to be logged.
<i>ap</i>	Provides a list of variables that is used to print additional information about the status of the logging process. The first argument needs to be a character string that describes what failed. The following arguments are expected in a format similar to the <b>printf</b> operation, which is a string format with the variables needed to fill the format variable places.
<i>va_alist</i>	Provides a variable list of parameters that includes the <i>lp</i> , <i>event</i> , and <i>ap</i> variables.

## Return Values

The **ll\_dbinit** and **ll\_hdinit** subroutines have no return values. The **\_ll\_log** and **ll\_log** subroutines return **OK** on success and **NOTOK** on failure.

## Related Information

The **isodetailor** subroutine.

Examples of SMUX Error Logging Routines, and SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## o\_number, o\_integer, o\_string, o\_igeneric, o\_generic, o\_specific, or o\_ipaddr Subroutine

### Purpose

Encodes values retrieved from the Management Information Base (MIB) into the specified variable binding.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/objects.h>
#include <isode/pepsy/SNMP-types.h>
#include <sys/types.h>
#include <netinet/in.h>
```

```

int o_number ( oi, v, number)
OI oi;
register struct type_SNMP_VarBind *v;
int number;
#define o_integer (oi, v, number) o_number ((oi), (v), (number))

int o_string (oi, v, base, len)
OI oi;
register struct type_SNMP_VarBind *v;
char *base;
int len;

int o_igeneric (oi, v, offset)
OI oi;
register struct type_SNMP_VarBind *v;
int offset;
int o_generic (oi, v, offset)
OI oi;
register struct type_SNMP_VarBind *v;
int offset;

int o_specific (oi, v, value)
OI oi;
register struct type_SNMP_VarBind *v;
caddr_t value;

int o_ipaddr (oi, v, netaddr)
OI oi;
register struct type_SNMP_VarBind *v;
struct sockaddr_in *netaddr;

```

## Description

The **o\_number** subroutine assigns a number retrieved from the MIB to the variable binding used to request it. Once an MIB value has been retrieved, the value must be stored in the binding structure associated with the variable requested. The **o\_number** subroutine places the integer *number* into the *v* parameter, which designates the binding for the variable. The *value* parameter type is defined by the *oi* parameter and is used to specify the encoding subroutine that stores the value. The *oi* parameter references a specific MIB variable and should be the same as the variable specified in the *v* parameter. The encoding functions are defined for each type of variable and are contained in the object identifier (**OI**) structure.

The **o\_integer** macro is defined in the `/usr/include/snmp/objects.h` file. This macro casts the *number* parameter as an integer. Use the **o\_integer** macro for types that are not integers but have integer values.

The **o\_string** subroutine assigns a string that has been retrieved for a MIB variable to the variable binding used to request the string. Once a MIB variable has been retrieved, the value is stored in the binding structure associated with the variable requested. The **o\_string** subroutine places the string, specified with the *base* parameter, into the variable binding in the *v* parameter. The length of the string represented in the *base* parameter equals the value of the *len* parameter. The length is used to define how much of the string is copied in the binding parameter of the variable. The *value* parameter type is defined by the *oi* parameter and is used to specify the encoding subroutine that stores the value. The *oi* parameter references a specific MIB variable and should be the same as the variable specified in the *v* parameter. The encoding subroutines are defined for each type of variable and are contained in the **OI** structure.

The **o\_generic** and **o\_igeneric** subroutines assign results that are already in the customer's MIB database. These two subroutines do not retrieve values from any other source. These subroutines check

whether the MIB database has information on how and what to encode as the value. The **o\_generic** and **o\_igeneric** subroutines also ensure that the variable requested is an instance. If the variable is an instance, the subroutines encode the value and return **OK**. The subroutine has an added set of return codes. If there is not any information about the variable, the subroutine returns **NOTOK** on a **get\_next** request and **int\_SNMP\_error\_\_status\_noSuchName** for the get and set requests. The difference between the **o\_generic** and the **o\_igeneric** subroutine is that the **o\_igeneric** subroutine provides a method for users to define a generic subroutine.

The **o\_specific** subroutine sets the binding value for a MIB variable with the value in a character pointer. The **o\_specific** subroutine ensures that the data-encoding procedure is defined. The encode subroutine is always checked by all of the **o\_** subroutines. The **o\_specific** subroutine returns the normal values.

The **o\_ipaddr** subroutine sets the binding value for variables that are network addresses. The **o\_ipaddr** subroutine uses the **sin\_addr** field of the **sockaddr\_in** structure to get the address. The subroutine does the normal checking and returns the results like the rest of the subroutines.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>oi</i>	Contains the <b>OI</b> data structure for the variable whose value is to be recorded into the binding structure.
<i>v</i>	Specifies the variable binding parameter, which is of type <b>type_SNMP_VarBind</b> . The <i>v</i> parameter contains a name and a value field. The value field contents are supplied by the <b>o_</b> subroutines.
<i>number</i>	Contains an integer to store in the value field of the <i>v</i> (variable bind) parameter.
<i>base</i>	Points to the character string to store in the value field of the <i>v</i> parameter.
<i>len</i>	Designates the length of the integer character string to copy. The character string is described by the <i>base</i> parameter.
<i>offset</i>	Contains an integer value of the current type of request, for example: type_SNMP_PDUs_get__next__request
<i>value</i>	Contains a character pointer to a value.
<i>netaddr</i>	Points to a <b>sockaddr_in</b> structure. The subroutine only uses the <b>sin_addr</b> field of this structure.

## Return Values

The return values for these subroutines are:

Value	Description
<b>int_SNMP_error__status_genErr</b>	Indicates an error occurred when setting the <i>v</i> parameter value.
<b>int_SNMP_error__status_noErr</b>	Indicates no errors found.

## Related Information

List of Network Manager Programming References.

SNMP Overview for Programmers, and Working with Management Information Base (MIB) Variables in *AIX Version 6.1 Communications Programming Concepts*.

---

## oid\_cmp, oid\_cpy, oid\_free, sprintoid, str2oid, ode2oid, oid2ode, oid2ode\_aux, prim2oid, or oid2prim Subroutine

### Purpose

Manipulates the object identifier data structure.

## Library

ISODE Library (**libisode.a**)

## Syntax

```
#include <isode/psap.h>
```

```
int oid_cmp (p, q)
OID p, q;
```

```
OID oid_cpy (oid)
OID oid;
```

```
void oid_free (oid)
OID oid;
```

```
char *sprintoid (oid)
OID oid;
```

```
OID str2oid (s)
char * s;
```

```
OID ode2oid (descriptor)
char * descriptor;
char *oid2ode (oid)
OID oid;
```

```
OID *oid2ode_aux (descriptor, quote)
char *descriptor;
int quote;
```

```
OID prim2oid (pe)
PE pe;
PE oid2prim (oid)
OID oid;
```

## Description

These subroutines are used to manipulate and translate object identifiers. The object identifier data (**OID**) structure and these subroutines are defined in the `/usr/include/isode/psap.h` file.

The **oid\_cmp** subroutine compares two **OID** structures. The **oid\_cpy** subroutine copies the object identifier, specified by the *oid* parameter, into a new structure. The **oid\_free** procedure frees the object identifier and does not have any return parameters.

The **sprintoid** subroutine takes an object identifier and returns the dot-notation description as a string. The string is in static storage and must be copied to other user storage if it is to be maintained. The **sprintoid** subroutine takes the object data and converts it without checking for the existence of the *oid* parameter.

The **str2oid** subroutine takes a character string specifying an object identifier in dot notation (for example, 1.2.3.6.1.2) and converts it into an **OID** structure. The space is static. To get a permanent copy of the **OID** structure, use the **oid\_cpy** subroutine.

The **oid2ode** subroutine is identical to the **sprintoid** subroutine except that the **oid2ode** subroutine checks whether the *oid* parameter is in the **isobjects** database. The **oid2ode** subroutine is implemented as a macro call to the **oid2ode\_aux** subroutine. The **oid2ode\_aux** subroutine is similar to the **oid2ode** subroutine except for an additional integer parameter that specifies whether the string should be enclosed by quotes. The **oid2ode** subroutine always encloses the string in quotes.

The **ode2oid** subroutine retrieves an object identifier from the **isobjects** database.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>p</i>	Specifies an <b>OID</b> structure.
<i>q</i>	Specifies an <b>OID</b> structure.
<i>descriptor</i>	Contains the object identifier descriptor data.
<i>oid</i>	Contains the object identifier data.
<i>s</i>	Contains a character string that defines an object identifier in dot notation.
<i>descriptor</i>	Contains the object identifier descriptor data.
<i>quote</i>	Specifies an integer that indicates whether a string should be enclosed in quotes. A value of 1 adds quotes; a value of 0 does not add quotes.
<i>pe</i>	Contains a presentation element in which the <b>OID</b> structure is encoded (as with the <b>oid2prim</b> subroutine) or decoded (as with the <b>prim2oid</b> subroutine).

## Return Values

The **oid\_cmp** subroutine returns a 0 if the structures are identical, -1 if the first object is less than the second, and a 1 if any other conditions are found. The **oid\_cpy** subroutine returns a pointer to the designated object identifier when the subroutine is successful.

The **oid2ode** subroutine returns the dot-notation description as a string in quotes. The **sprintoid** subroutine returns the dot-notation description as a string without quotes.

The **ode2oid** subroutine returns a static pointer to the object identifier. If the **ode2oid** and **oid\_cpy** subroutines are not successful, the **NULLOID** value is returned.

## Related Information

The **oid\_extend** subroutine, **oid\_normalize** subroutine.

List of Network Manager Programming References.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## oid\_extend or oid\_normalize Subroutine

### Purpose

Extends the base ISODE library subroutines.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/objects.h>
```

```
OID oid_extend (q, howmuch)
```

```
OID q;
```

```
integer howmuch;
```

```
OID oid_normalize (q, howmuch, initial)
OID q;
integer howmuch, initial;
```

## Description

The **oid\_extend** subroutine is used to extend the current object identifier data (**OID**) structure. The **OID** structure contains an integer number of entries and an array of integers. The **oid\_extend** subroutine creates a new, extended **OID** structure with an array of the size specified in the *howmuch* parameter plus the original array size specified in the *q* parameter. The original values are copied into the first entries of the new structure. The new values are uninitialized. The entries of the **OID** structure are used to represent the values of an Management Information Base (MIB) tree in dot notation. Each entry represents a level in the MIB tree.

The **oid\_normalize** subroutine extends and adjusts the values of the **OID** structure entries. The **oid\_normalize** subroutine extends the **OID** structure and then decrements all nonzero values by 1. The new values are initialized to the value of the *initial* parameter. This subroutine stores network address and netmask information in the **OID** structure.

These subroutines do not free the *q* parameter.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>q</i>	Specifies the size of the original array.
<i>howmuch</i>	Specifies the size of the array for the new <b>OID</b> structure.
<i>initial</i>	Indicates the initialized value of the <b>OID</b> structure extensions.

## Return Values

Both subroutines, when successful, return the pointer to the new object identifier structure. If the subroutines fail, the **NULLOID** value is returned.

## Related Information

The **oid\_cmp**, **oid\_cpy**, **oid\_free**, **sprintoid**, **str2oid**, **ode2oid**, **oid2ode**, **oid2ode\_aux**, **prim2oid**, or **oid2prim** subroutine.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## readobjects Subroutine

### Purpose

Allows the SNMP multiplexing (SMUX) peer to read the Management Information Base (MIB) variable structure.

### Library

SNMP Library (**libsnp.a**)

### Syntax

```
#include <isode/snp/objects.h>
```

```
int
readobjects ( file)
char *file;
```

## Description

The **readobjects** subroutine reads the file given in the *file* parameter. This file must contain the MIB variable descriptions that the SMUX peer supports. The SNMP library functions require basic information about the MIB tree supported by the SMUX peer. These structures are supplied from information in the **readobjects** file. The **text2oid** subroutine receives a string description and uses the object identifier information retrieved with the **readobjects** subroutine to return a MIB object identifier. The file designated in the *file* parameter must be in the following form:

```
<MIB directory> <MIB position>
```

```
<MIB name> <MIB position> <MIB type> <MIB access> <MIB required?>
```

```
<MIB name> <MIB position> <MIB type> <MIB access> <MIB required?>
```

```
...
```

An example of a file that uses this format is **/etc/mib.defs**. The **/etc/mib.defs** file defines the MIBII tree used in the SNMP agent.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

*file*        Contains the name of the file to be read. If the value is NULL, the **/etc/mib.defs** file is read.

## Return Values

If the subroutine is successful, **OK** is returned. Otherwise, **NOTOK** is returned.

## Related Information

The **text2oid** subroutine.

RFC 1155 describes the basic MIB structure.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## s\_generic Subroutine

### Purpose

Sets the value of the Management Information Base (MIB) variable in the database.

### Library

The SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/objects.h>
```

```
int s_generic
```

```
( oi, v, offset)
```

```
OI oi;
```

```
register struct type_SNMP_VarBind *v;
```

```
int offset;
```

## Description

The **s\_generic** subroutine sets the database value of the MIB variable. The subroutine retrieves the information it needs from a value in a variable binding within the Protocol Data Unit (PDU). The **s\_generic** subroutine sets the MIB variable, specified by the object identifier *oi* parameter, to the value field specified by the *v* parameter.

The *offset* parameter is used to determine the stage of the set process. If the *offset* parameter value is **type\_SNMP\_PDUs\_set\_\_request**, the value is checked for validity and the value in the *ot\_save* field in the **OI** structure is set. If the *offset* parameter value is **type\_SNMP\_PDUs\_commit**, the value in the *ot\_save* field is freed and moved to the MIB *ot\_info* field. If the *offset* parameter value is **type\_SNMP\_PDUs\_rollback**, the value in the *ot\_save* field is freed and no new value is written.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>oi</i>	Designates the <b>OI</b> structure representing the MIB variable to be set.
<i>v</i>	Specifies the variable binding that contains the value to be set.
<i>offset</i>	Contains the stage of the set. The possible values for the <i>offset</i> parameter are <b>type_SNMP_PDUs_commit</b> , <b>type_SNMP_PDUs_rollback</b> , or <b>type_SNMP_PDUs_set__request</b> .

## Return Values

If the subroutine is successful, a value of **int\_SNMP\_error\_\_status\_noError** is returned. Otherwise, a value of **int\_SNMP\_error\_\_status\_badValue** is returned.

## Related Information

The **o\_number**, **o\_integer**, **o\_string**, **o\_specific**, **o\_igeneric**, **o\_generic**, or **o\_ipaddr** subroutines.

SNMP Overview for Programmers, and SNMP daemon processing in *Networks and communication management*.

---

## smux\_close Subroutine

### Purpose

Ends communications with the SNMP agent.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_close ( reason)
int reason;
```

### Description

The **smux\_close** subroutine closes the transmission control protocol (TCP) connection from the SNMP multiplexing (SMUX) peer. The **smux\_close** subroutine sends the close protocol data unit (PDU) with the error code set to the *reason* value. The subroutine closes the TCP connection and frees the socket. This subroutine also frees information it was maintaining for the connection.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

*reason* Indicates an integer value denoting the reason the close PDU message is being sent.

## Return Values

If the subroutine is successful, **OK** is returned. Otherwise, **NOTOK** is returned.

## Error Codes

If the subroutine returns **NOTOK**, the **smux\_errno** global variable is set to one of the following values:

Value	Description
<b>invalidOperation</b>	Indicates that the <b>smux_init</b> subroutine has not been executed successfully.
<b>congestion</b>	Indicates that memory could not be allocated for the close PDU. The TCP connection is closed.
<b>youLoseBig</b>	Indicates that the SNMP code has a problem. The TCP connection is closed.

## Related Information

The **smux\_error** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_error Subroutine

### Purpose

Creates a readable string from the **smux\_errno** global variable value.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>

char *smux_error ( error)
int error;
```

### Description

The **smux\_error** subroutine creates a readable string from error code values in the **smux\_errno** global variable in the **smux.h** file. The **smux** global variable, **smux\_errno**, is set when an error occurs. The **smux\_error** subroutine can also get a string that interprets the value of the **smux\_errno** variable. The **smux\_error** subroutine can be used to retrieve any numbers, but is most useful interpreting the integers returned in the **smux\_errno** variable.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

*error*       Contains the error to interpret. Usually called with the value of the **smux\_errno** variable, but can be called with any error that is an integer.

## Return Values

If the subroutine is successful, a pointer to a static string is returned. If an error occurs, a string of the type SMUX error %s(%d) is returned. The %s value is a string representing the explanation of the error. The %d is the number used to reference that error.

## Related Information

The **smux\_close** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_free\_tree Subroutine

### Purpose

Frees the object tree when a **smux** tree is unregistered.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
void smux_free_tree ( parent, child)
char *parent;
char *child;
```

### Description

The **smux\_free\_tree** subroutine frees elements in the Management Information Base (MIB) list within an SNMP multiplexing (SMUX) peer. If the SMUX peer implements the MIB list with the **readobjects** subroutine, a list of MIBs is created and maintained. These MIBs are kept in the object tree (**OT**) data structures.

Unlike the **smux\_register** subroutine, the **smux\_free\_tree** subroutine frees the MIB elements even if the tree is unregistered by the **snmpd** daemon. This functionality is not performed by the **smux\_register** routine because the **OT** list is created independently of registering a tree with the **snmpd** daemon. The unregistered objects should be removed as the user deems appropriate. Remove the unregistered objects if the **smux** peer is highly dynamic. If the peer registers and unregisters many trees, it might be reasonable to add and delete the **OT** MIB list on the fly. The **smux\_free\_tree** subroutine expects the parent of the MIB tree in the local **OT** list to delete unregistered objects.

This subroutine does not return values or error codes.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>parent</i>	Contains a character string holding the immediate parent of the tree to be deleted.
<i>child</i>	Contains a character string holding the beginning of the tree to be deleted.

The character strings are names or dot notations representing object identifiers.

## Related Information

The **snmpd** command.

The **readobjects** subroutine, **smux\_register** subroutine.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_init Subroutine

### Purpose

Initiates the transmission control protocol (TCP) socket that the SNMP multiplexing (SMUX) agent uses and clears the basic SMUX data structures.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_init ( debug)
int debug;
```

### Description

The **smux\_init** subroutine initializes the TCP socket used by the SMUX agent to talk to the SNMP daemon. The subroutine assumes that loopback will be used to define the path to the SNMP daemon. The subroutine also clears the base structures the SMUX code uses. This subroutine also sets the debug level that is used when running the SMUX subroutines.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

### Parameters

*debug*      Indicates the level of debug to be printed during SMUX subroutines.

### Return Values

If the subroutine is successful, the socket descriptor is returned. Otherwise, the value of **NOTOK** is returned and the **smux\_errno** global variable is set.

## Error Codes

Possible values for the **smux\_errno** global variable are:

Value	Description
<b>congestion</b>	Indicates memory allocation problems
<b>youLoseBig</b>	Signifies problem with SNMP library code
<b>systemError</b>	Indicates TCP connection failure.

These are defined in the **/usr/include/isode/snmp/smux.h** file.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_register Subroutine

### Purpose

Registers a section of the Management Information Base (MIB) tree with the Simple Network Management Protocol (SNMP) agent.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_register ( subtree, priority, operation)
```

```
OID subtree;  
int priority;  
int operation;
```

### Description

The **smux\_register** subroutine registers the section of the MIB tree for which the SMUX peer is responsible with the SNMP agent. Using the **smux\_register** subroutine, the SMUX peer informs the SNMP agent of both the level of responsibility the SMUX peer has and the sections of the MIB tree for which it is responsible. The level of responsibility (priority) the SMUX peer sends determines which requests it can answer. Lower priority numbers correspond to higher priority.

If a tree is registered more than once, the SNMP agent sends requests to the registered SMUX peer with the highest priority. If the priority is set to -1, the SNMP agent attempts to give the SMUX peer the highest available priority. The *operation* parameter defines whether the MIB tree is added with readOnly or readWrite permissions, or if it should be deleted from the list of register trees. The SNMP agent returns an acknowledgment of the registration. The acknowledgment indicates the success of the registration and the actual priority received.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>subtree</i>	Indicates an object identifier that contains the root of the MIB tree to be registered.
<i>priority</i>	Indicates the level of responsibility that the SMUX peer has on the MIB tree. The priority levels range from 0 to ( $2^{31} - 2$ ). The lower the priority number, the higher the priority. A priority of -1 tells the SNMP daemon to assign the highest priority currently available.
<i>operation</i>	Specifies the operation for the SNMP agent to apply to the MIB tree. Possible values are <b>delete</b> , <b>readOnly</b> , or <b>readWrite</b> . The <b>delete</b> operation removes the MIB tree from the SMUX peers in the eyes of the SNMP agent. The other two values specify the operations allowed by the SMUX peer on the MIB tree that is being registered with the SNMP agent.

## Return Values

The values returned by this subroutine are **OK** on success and **NOTOK** on failure.

## Error Codes

If the subroutine is unsuccessful, the **smux\_errno** global variable is set to one of the following values:

Value	Description
<b>parameterMissing</b>	Indicates a parameter was null. When the parameter is fixed, the <b>smux_register</b> subroutine can be reissued.
<b>invalidOperation</b>	Indicates that the <b>smux_register</b> subroutine is trying to perform this operation before a <b>smux_init</b> operation has successfully completed. Start over with a new <b>smux_init</b> subroutine call.
<b>congestion</b>	Indicates a memory problem occurred. The TCP connection is closed. Start over with a new <b>smux_init</b> subroutine call.
<b>youLoseBig</b>	Indicates an SNMP code problem has occurred. The TCP connection is closed. Start over with a new <b>smux_init</b> subroutine call.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_init** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

RFC1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_response Subroutine

### Purpose

Sends a response to a Simple Network Management Protocol (SNMP) agent.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_response ( event )  
struct type_SNMP_GetResponse__PDU *event;
```

## Description

The **smux\_response** subroutine sends a protocol data unit (PDU), also called an event, to the SNMP agent. The subroutine does not check whether the Management Information Base (MIB) tree is properly registered. The subroutine checks only to see whether a Transmission Control Protocol (TCP) connection to the SNMP agent exists and ensures that the *event* parameter is not null.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

*event* Specifies a **type SNMP\_GetResponse\_\_PDU** variable that contains the response PDU to send to the SNMP agent.

## Return Values

If the subroutine is successful, **OK** is returned. Otherwise, **NOTOK** is returned.

## Error Codes

If the subroutine is unsuccessful, the **smux\_errno** global variable is set to one of the following values:

Value	Description
<b>parameterMissing</b>	Indicates the parameter was null. When the parameter is fixed, the subroutine can be reissued.
<b>invalidOperation</b>	Indicates the subroutine was attempted before the <b>smux_init</b> subroutine successfully completed. Start over with the <b>smux_init</b> subroutine.
<b>youLoseBig</b>	Indicates a SNMP code problem has occurred and the TCP connection is closed. Start over with the <b>smux_init</b> subroutine.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_simple\_open Subroutine

### Purpose

Sends the open protocol data unit (PDU) to the Simple Network Management Protocol (SNMP) daemon.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_simple_open (identity, description, commname, commlen)
OID identity;
char * description;
char * commname;
int commlen;
```

## Description

Following the **smux\_init** command, the **smux\_simple\_open** subroutine alerts the SNMP daemon that incoming messages are expected. Communication with the SNMP daemon is accomplished by sending an open PDU to the SNMP daemon. The **smux\_simple\_open** subroutine uses the *identity* object-identifier parameter to identify the SNMP multiplexing (SMUX) peer that is starting to communicate. The *description* parameter describes the SMUX peer. The *commname* and the *commlen* parameters supply the password portion of the open PDU. The *commname* parameter is the password used to authenticate the SMUX peer. The SNMP daemon finds the password in the **/etc/snmpd.conf** file. The SMUX peer can store the password in the **/etc/snmpd.peers** file. The *commlen* parameter specifies the length of the *commname* parameter value.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>identity</i>	Specifies an object identifier that describes the SMUX peer.
<i>description</i>	Contains a string of characters that describes the SMUX peer. The <i>description</i> parameter value cannot be longer than 254 characters.
<i>commname</i>	Contains the password to be sent to the SNMP agent. Can be a null value.
<i>commlen</i>	Indicates the length of the community name ( <i>commname</i> parameter) to be sent to the SNMP agent. The value for this parameter must be at least 0.

## Return Values

The subroutine returns an integer value of **OK** on success or **NOTOK** on failure.

## Error Codes

If the subroutine is unsuccessful, the **smux\_errno** global variable is set one of the following values:

Value	Description
<b>parameterMissing</b>	Indicates that a parameter was null. The <i>commname</i> parameter can be null, but the <i>commlen</i> parameter value should be at least 0.
<b>invalidOperation</b>	Indicates that the <b>smux_init</b> subroutine did not complete successfully before the <b>smux_simple_open</b> subroutine was attempted. Correct the parameters and reissue the <b>smux_simple_open</b> subroutine.
<b>inProgress</b>	Indicates that the <b>smux_init</b> call has not completed the TCP connection. The <b>smux_simple_open</b> can be reissued.
<b>systemError</b>	Indicates the TCP connection was not completed. Do not reissue this subroutine without restarting the process with a <b>smux_init</b> subroutine call.
<b>congestion</b>	Indicates a lack of available memory space. Do not reissue this subroutine without restarting the process with a <b>smux_init</b> subroutine call.
<b>youLoseBig</b>	The SNMP code is having problems. Do not reissue this subroutine without restarting the process with a <b>smux_init</b> subroutine call.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_trap** subroutine, **smux\_wait** subroutine.

List of Network Manager Programming References.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_trap Subroutine

### Purpose

Sends SNMP multiplexing (SMUX) peer traps to the Simple Network Management Protocol (SNMP) agent.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_trap ( generic, specific, bindings)
int generic;
int specific;
struct type_SNMP_VarBindList *bindings;
```

### Description

The **smux\_trap** subroutine allows the SMUX peer to generate traps and send them to the SNMP agent. The subroutine sets the *generic* and *specific* fields in the trap packet to values specified by the parameters. The subroutine also allows the SMUX peer to send a list of variable bindings to the SNMP agent. The variable bindings are values associated with specific variables. If the trap is to return a set of variables, the variables are sent in the variable binding list.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

### Parameters

<i>generic</i>	Contains an integer specifying the generic trap type. The value must be one of the following:
0	Specifies a cold start.
1	Specifies a warm start.
2	Specifies a link down.
3	Specifies a link up.
4	Specifies an authentication failure.
5	Specifies an EGP neighbor loss.
6	Specifies an enterprise-specific trap type.
<i>specific</i>	Contains an integer that uniquely identifies the trap. The unique identity is typically assigned by the registration authority for the enterprise owning the SMUX peer.
<i>bindings</i>	Indicates the variable bindings to assign to the trap protocol data unit (PDU).

### Return Values

The subroutine returns **NOTOK** on failure and **OK** on success.

### Error Codes

If the subroutine is unsuccessful, the **smux\_errno** global variable is set to one of the following values:

Value	Description
<b>invalidOperation</b>	Indicates the Transmission Control Protocol (TCP) connection was not completed.
<b>congestion</b>	Indicates memory is not available. The TCP connection was closed.
<b>youLoseBig</b>	Indicates an error occurred in the SNMP code. The TCP connection was closed.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_wait** subroutine.

RFC 1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## smux\_wait Subroutine

### Purpose

Waits for a message from the Simple Network Management Protocol (SNMP) agent.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/smux.h>
```

```
int smux_wait ( event,  isecs)
struct type_SMUX_PDUs **event;
int isecs;
```

### Description

The **smux\_wait** subroutine waits for a period of seconds, designated by the value of the *isecs* parameter, and returns the protocol data unit (PDU) received. The **smux\_wait** subroutine waits on the socket descriptor that is initialized in a **smux\_init** subroutine and maintained in the SMUX subroutines. The **smux\_wait** subroutine waits up to *isecs* seconds. If the value of the *isecs* parameter is 0, the **smux\_wait** subroutine returns only the first packet received. If the value of the *isecs* parameter is less than 0, the **smux\_wait** subroutine waits indefinitely for the next message or returns a message already received. If no data is received, the **smux\_wait** subroutine returns an error message of **NOTOK** and sets the **smux\_errno** variable to the **inProgress** value. If the **smux\_wait** subroutine is successful, it returns the first PDU waiting to be received. If a close PDU is received, the subroutine will automatically close the TCP connection and return **OK**.

This subroutine is part of the SNMP Application Programming Interface in the TCP/IP facility.

### Parameters

<i>event</i>	Points to a pointer of <b>type_SMUX_PDUs</b> . This holds the PDUs received by the <b>smux_wait</b> subroutine.
<i>isecs</i>	Specifies an integer value equal to the number of seconds to wait for a message.

### Return Values

If the subroutine is successful, the value **OK** is returned. Otherwise, the return value is **NOTOK**.

## Error Codes

If the subroutine is unsuccessful, the **smux\_errno** global variable is set to one of the following values:

Value	Description
<b>parameterMissing</b>	Indicates that the <i>event</i> parameter value was null.
<b>inProgress</b>	Indicates that there was nothing for the subroutine to receive.
<b>invalidOperation</b>	Indicates that the <b>smux_init</b> subroutine was not called or failed to operate.
<b>youLoseBig</b>	Indicates an error occurred in the SNMP code. The TCP connection was closed.

## Related Information

The **smux\_close** subroutine, **smux\_error** subroutine, **smux\_init** subroutine, **smux\_register** subroutine, **smux\_response** subroutine, **smux\_simple\_open** subroutine, **smux\_trap** subroutine.

RFC1227, SNMP MUX Protocol and MIB.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## text2inst, name2inst, next2inst, or nexttot2inst Subroutine

### Purpose

Retrieves instances of variables from various forms of data.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/objects.h>
```

```
0I text2inst ( text)
char *text;
```

```
0I name2inst ( oid)
0ID oid;
0I next2inst (oid)
0ID oid;
```

```
0I nexttot2inst (oid, ot)
0ID oid;
0T ot;
```

### Description

These subroutines return pointers to the actual objects in the database. When supplied with a way to identify the object, the subroutines return the corresponding object.

The **text2inst** subroutine takes a character string object identifier from the *text* parameter. The object's database is then examined for the specified object. If the specific object is not found, the **NULLOI** value is returned.

The **name2inst** subroutine uses an object identifier structure specified in the *oid* parameter to specify which object is desired. If the object cannot be found, a **NULLOI** value is returned.

The **next2inst** and **nextot2inst** subroutines find the next object in the database given an object identifier. The **next2inst** subroutine starts at the root of the tree, while the **nextot2inst** subroutine starts at the object given in the *ot* parameter. If another object cannot be found, the **NULLOI** value will be returned.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>text</i>	Specifies the character string used to identify the object wanted in the <b>text2inst</b> subroutine.
<i>oid</i>	Specifies the object identifier structure used to identify the object wanted in the <b>name2inst</b> , <b>next2inst</b> , and <b>nextot2inst</b> subroutines.
<i>ot</i>	Specifies an object in the database used as a starting point for the <b>nextot2inst</b> subroutine.

## Return Values

If the subroutine is successful, an **OI** value is returned. **OI** is a pointer to an object in the database. On a failure, a **NULLOI** value is returned.

## Related Information

The **text2oid** subroutine, **text2obj** subroutine.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## text2oid or text2obj Subroutine

### Purpose

Converts a text string into some other value.

### Library

SNMP Library (**libsnmp.a**)

### Syntax

```
#include <isode/snmp/objects.h>
```

```
OID text2oid ( text)
```

```
char *text;
```

```
OT text2obj (text)
```

```
char *text;
```

### Description

The **text2oid** subroutine takes a character string and returns an object identifier. The string can be a name, a name.numbers, or dot notation. The returned object identifier is in memory-allocation storage and should be freed when the operation is completed with the **oid\_free** subroutine.

The **text2obj** subroutine takes a character string and returns an object. The string needs to be the name of a specific object. The subroutine returns a pointer to the object.

These subroutines are part of the SNMP Application Programming Interface in the TCP/IP facility.

## Parameters

<i>text</i>	Contains a text string used to specify the object identifier or object to be returned.
-------------	--

## Return Values

On a successful execution, these subroutines return completed data structures. If a failure occurs, the **text2oid** subroutine returns a **NULLOID** value and the **text2obj** returns a **NULLOT** value.

## Related Information

The **malloc** subroutine, **oid\_free** subroutine, **text2inst** subroutine.

SNMP Overview for Programmers in *AIX Version 6.1 Communications Programming Concepts*.

---

## Chapter 2. Sockets

---

### **`_getlong` Subroutine**

#### **Purpose**

Retrieves long byte quantities.

#### **Library**

Standard C Library (**libc.a**)

#### **Syntax**

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

```
unsigned long _getlong ( MessagePtr)
u_char *MessagePtr;
```

#### **Description**

The **`_getlong`** subroutine gets long quantities from the byte stream or arbitrary byte boundaries.

The **`_getlong`** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolves domain names. Global information used by the resolver subroutines is kept in the **`_res`** data structure. The **`/usr/include/resolv.h`** file contains the **`_res`** structure definition.

All applications containing the **`_getlong`** subroutine must be compiled with the **`_BSD`** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **`libbsd.a`** library.

#### **Parameters**

*MessagePtr*                Specifies a pointer into the byte stream.

#### **Return Values**

The **`_getlong`** subroutine returns an unsigned long (32-bit) value.

#### **Files**

**`/etc/resolv.conf`**                Lists name server and domain names.

#### **Related Information**

The **`dn_comp`** subroutine, **`dn_expand`** subroutine, **`_getshort`** subroutine, **`putlong`** subroutine, **`putshort`** subroutine, **`res_init`** subroutine, **`res_mkquery`** subroutine, “res\_ninit Subroutine” on page 176, **`res_query`** subroutine, **`res_search`** subroutine, **`res_send`** subroutine.

Sockets Overview, and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## **\_\_getshort Subroutine**

### **Purpose**

Retrieves short byte quantities.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

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

```
unsigned short getshort ( MessagePtr)
u_char *MessagePtr;
```

### **Description**

The **\_\_getshort** subroutine gets quantities from the byte stream or arbitrary byte boundaries.

The **\_\_getshort** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. Global information used by the resolver subroutines is kept in the **\_\_res** data structure. The **/usr/include/resolv.h** file contains the **\_\_res** structure definition.

All applications containing the **\_\_getshort** subroutine must be compiled with the **\_\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### **Parameters**

*MessagePtr*                Specifies a pointer into the byte stream.

### **Return Values**

The **\_\_getshort** subroutine returns an unsigned short (16-bit) value.

### **Files**

**/etc/resolv.conf**                                Defines name server and domain names.

### **Related Information**

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_\_getlong** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176 **res\_send** subroutine.

Sockets Overview, and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## **\_putlong Subroutine**

### **Purpose**

Places long byte quantities into the byte stream.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

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

```
void _putlong ( Long, MessagePtr)
unsigned long Long;
u_char *MessagePtr;
```

### **Description**

The **\_putlong** subroutine places long byte quantities into the byte stream or arbitrary byte boundaries.

The **\_putlong** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. Global information used by the resolver subroutines is kept in the **\_res** data structure. The **/usr/include/resolv.h** file contains the **\_res** structure definition.

All applications containing the **\_putlong** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### **Parameters**

<i>Long</i>	Represents a 32-bit integer.
<i>MessagePtr</i>	Represents a pointer into the byte stream.

### **Files**

<b>/etc/resolv.conf</b>	Lists the name server and domain name.
-------------------------	--

### **Related Information**

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_query** subroutine, **res\_search** subroutine, **res\_send** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## **\_putshort Subroutine**

### **Purpose**

Places short byte quantities into the byte stream.

### **Library**

Standard C Library (**libc.a**)

### **Syntax**

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

```
void _putshort ( Short, MessagePtr)
unsigned short Short;
u_char *MessagePtr;
```

### **Description**

The **\_putshort** subroutine puts short byte quantities into the byte stream or arbitrary byte boundaries.

The **\_putshort** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. Global information used by the resolver subroutines is kept in the **\_res** data structure. The **/usr/include/resolv.h** file contains the **\_res** structure definition.

All applications containing the **\_putshort** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### **Parameters**

<i>Short</i>	Represents a 16-bit integer.
<i>MessagePtr</i>	Represents a pointer into the byte stream.

### **Files**

<b>/etc/resolv.conf</b>	Lists the name server and domain name.
-------------------------	--

### **Related Information**

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_send** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## accept Subroutine

### Purpose

Accepts a connection on a socket to create a new socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int accept ( Socket, Address, AddressLength)
int Socket;
struct sockaddr *Address;
socklen_t *AddressLength;
```

### Description

The **accept** subroutine extracts the first connection on the queue of pending connections, creates a new socket with the same properties as the specified socket, and allocates a new file descriptor for that socket.

If the **listen** queue is empty of connection requests, the **accept** subroutine:

- Blocks a calling socket of the blocking type until a connection is present.
- Returns an **EWOULDBLOCK** error code for sockets marked nonblocking.

The accepted socket cannot accept more connections. The original socket remains open and can accept more connections.

The **accept** subroutine is used with **SOCK\_STREAM** and **SOCK\_CONN\_DGRAM** socket types.

For **SOCK\_CONN\_DGRAM** socket type and **ATM** protocol, a socket is not ready to transmit/receive data until **SO\_ATM\_ACCEPT** socket option is called. This allows notification of an incoming connection to the application, followed by modification of appropriate parameters and then indicate that a connection can become fully operational.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

### Parameters

<i>Socket</i>	Specifies a socket created with the <b>socket</b> subroutine that is bound to an address with the <b>bind</b> subroutine and has issued a successful call to the <b>listen</b> subroutine.
<i>Address</i>	Specifies a result parameter that is filled in with the address of the connecting entity as known to the communications layer. The exact format of the <i>Address</i> parameter is determined by the domain in which the communication occurs.
<i>AddressLength</i>	Specifies a parameter that initially contains the amount of space pointed to by the <i>Address</i> parameter. Upon return, the parameter contains the actual length (in bytes) of the address returned. The <b>accept</b> subroutine is used with <b>SOCK_STREAM</b> socket types.

### Return Values

Upon successful completion, the **accept** subroutine returns the nonnegative socket descriptor of the accepted socket.

If the **accept** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **accept** subroutine is unsuccessful if one or more of the following is true:

<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	The connection has been reset by the partner.
<b>EINTR</b>	The <b>accept</b> function was interrupted by a signal that was caught before a valid connection arrived.
<b>EINVAL</b>	The socket referenced by <i>s</i> is not currently a listen socket or has been shutdown with <b>shutdown</b> . A <b>listen</b> must be done before an <b>accept</b> is allowed.
<b>EMFILE</b>	The system limit for open file descriptors per process has already been reached ( <b>OPEN_MAX</b> ).
<b>ENFILE</b>	The maximum number of files allowed are currently open.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The referenced socket is not of type <b>SOCK_STREAM</b> .
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EWouldBlock</b>	The socket is marked as nonblocking, and no connections are present to be accepted.
<b>ENETDOWN</b>	The network with which the socket is associated is down.
<b>ENOTCONN</b>	The socket is not in the connected state.
<b>ECONNABORTED</b>	The client aborted the connection.

## Examples

As illustrated in this program fragment, once a socket is marked as listening, a server process can accept a connection:

```
struct sockaddr_in from;  
.  
.  
.  
fromlen = sizeof(from);  
newsock = accept(socket, (struct sockaddr*)&from, &fromlen);
```

## Related Information

The **connect** subroutine, **getsockname** subroutine, **listen** subroutine, **socket** subroutine.

Accepting UNIX Stream Connections Example Program, Binding Names to Sockets, Sockets Overview, Understanding Socket Connections, and Understanding Socket Creation in *AIX Version 6.1 Communications Programming Concepts*.

---

## arpresolve\_common Subroutine

### Purpose

Reads or creates new arp entries so that hardware addresses can be resolved.

### Syntax

```
int arpresolve_common ( ac, m, arpwhohas, dst, hwaddr, szhwaddr, extra, if_dependent)  
  
    register struct arpcom *ac;  
    struct mbuf *m;  
    int (*arpwhohas)(register struct arpcom *ac,  
                     struct in_addr *addr, int skipbestif, void *extra),  
    struct sockaddr_in *dst;
```

```

u_char * hwaddr;
int szhwaddr;
void *extra;
union if_dependent *if_dependent;

```

## Description

The **arpresolve\_common** subroutine reads or creates new arp entries so that hardware addresses can be resolved. It is called by **arpresolve** from the IF layer of the interface. If the arp entry is complete, then **arpresolve\_common** returns the address pointed to by *hwaddr* and the data pointed to by *if\_dependent* if *if\_dependent* is true. If the arp entry is not complete, then this subroutine adds the memory buffer pointed to by *mbuf* to **at\_hold**. **at\_hold** holds one or more packets that are waiting for the arp entry to complete so they can be transmitted.

If an arp entry does not exist, **arpresolve\_common** creates a new entry by calling **arptnew** and then adds the memory buffer pointed to by *mbuf* to **at\_hold**. This subroutine calls **arpwhoahas** when it creates a new arp entry or when the timer for the incomplete arp entry (with the IP address that is pointed to by *dst*) has expired.

## Parameters

<i>ac</i>	Points to the <b>arpcom</b> structure.
<i>m</i>	Points to the memory buffer ( <b>mbuf</b> ), which will be added to the list awaiting completion of the arp table entry.
<i>arpwhoahas</i>	Points to the <b>arpwhoahas</b> subroutine.
<i>addr</i>	Points to the <b>in_addr</b> structure's address.
<i>extra</i>	A void pointer that can be used in the future so that IF layers can pass extra structures to <b>arpwhoahas</b> .
<i>dst</i>	Points to the <b>sockaddr_in</b> structure. This structure has the destination IP address.
<i>hwaddr</i>	Points to the buffer. This buffer contains the hardware address if it finds a completed entry.
<i>szhwaddr</i>	Size of the buffer pointed to by <i>hwaddr</i> .
<i>if_dependent</i>	Pointer to the <b>if_dependent</b> structure. <b>arpresolve_common</b> uses this to pass the <b>if_dependent</b> data, which is part of the <b>arptab</b> entry, to the calling function.

## Return Values

<b>ARP_MBUF</b>	The arp entry is not complete.
<b>ARP_HWADDR</b>	The <i>hwaddr</i> buffer is filled with the hardware address.
<b>ARP_FLG_NOARP</b>	The arp entry does not exist, and the <b>IFF_NOARP</b> flag is set only if the value of <b>if_type</b> is <b>IFT_ETHER</b> .

## Related Information

“arpupdate Subroutine”

---

## arpupdate Subroutine

### Purpose

Updates arp entries for a given IP address.

## Syntax

```
int arpupdate (ac, m, hp, action, prm)
    register struct arpcom *ac;
    struct mbuf *m;
    caddr_t hp;
    int action;
    struct arpupdate_parm *prm;
```

## Description

The **arpupdate** subroutine updates arp entries for a given IP address. It is called by **arpinput** from the IF layer of the interface. This subroutine searches the arp table for an entry that matches the IP address. It then updates the arp entry for the given IP address. The **arpupdate** subroutine also performs reverse arp lookups.

The **arpupdate** subroutine enters a new address in **arptab**, pushing out the oldest entry from the bucket if there is no room. This subroutine always succeeds because no bucket can be completely filled with permanent entries (except when **arpiocli** tests whether another permanent entry can fit).

Depending on the action specified, the prm IP addresses **isaddr**, **itaddr**, and **myaddr** are used by the **arpupdate** subroutine.

## Parameters

<i>ac</i>	Points to the <b>arpcom</b> structure.
<i>m</i>	Points to the memory buffer ( <b>mbuf</b> ), that contains the arp response packet received by the interface.
<i>hp</i>	Points to the buffer that is passed by the interrupt handler.
<i>action</i>	Returns a value that indicates which action is taken:  <b>LOOK</b> Looks for the <b>isaddr</b> IP address in the arp table and returns the hardware address and <b>if_dependent</b> structure.  <b>LKPUB</b> Looks for the <b>isaddr</b> IP address in the arp table and returns the hardware address and <b>if_dependent</b> structure only if the <b>ATF_PUBL</b> is set.  <b>UPDT</b> Updates the arp entry for an IP address ( <b>isaddr</b> ). If no arp entry is there, creates a new one and updates the <b>if_dependent</b> structure using the <b>ptr</b> function passed in the <b>prm</b> structure.  <b>REVARP</b> Reverses the arp request. <i>hwaddr</i> contains the hardware address, <i>szhwaddr</i> indicates its size, and <i>saddr</i> returns the IP address if an entry is found.
<i>prm</i>	Points to the <b>arpupdate_parm</b> structure. The values are:  <b>LOOK or LKPUB</b> <b>itaddr</b> and <b>myaddr</b> are ignored. <b>isaddr</b> is used for arp table lookup.  <b>UPDTE</b> <b>isaddr</b> points to the sender protocol address. <b>itaddr</b> points to the target protocol address. <b>myaddr</b> points to the protocol address of the interface that received the packet.

## Return Values

<b>ARP_OK</b>	Lookup or update was successful.
<b>ARP_FAIL</b>	Lookup or update failed.
<b>ARP_NEWF</b>	New arp entry could not be created.

## Related Information

“arpresolve\_common Subroutine” on page 30

---

## bind Subroutine

### Purpose

Binds a name to a socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int bind ( Socket, Name, NameLength)
int Socket;
const struct sockaddr *Name;
socklen_t NameLength;
```

### Description

The **bind** subroutine assigns a *Name* parameter to an unnamed socket. Sockets created by the **socket** subroutine are unnamed; they are identified only by their address family. Subroutines that connect sockets either assign names or use unnamed sockets.

For a UNIX domain socket, a **connect** call only succeeds if the process that calls **connect** has read and write permissions on the socket file created by the **bind** call. Permissions are determined by the **umask** value of the process that created the file.

An application program can retrieve the assigned socket name with the **getsockname** subroutine.

The socket applications can be compiled with **COMPAT\_43** defined. This makes the **sockaddr** structure BSD 4.3 compatible. For more details refer to the **socket.h** file.

Binding a name in the UNIX domain creates a socket in the file system that must be deleted by the caller when it is no longer needed.

**Note:** When you enable IPv6 for an application, IPv4 addresses are also supported. You can use an AF\_INET6 socket to send and receive both IPv4 and IPv6 packets because AF\_INET6 sockets are capable of handling communication with both IPv4 and IPv6 hosts. However, you must convert the address format of the IPv4 addresses that were previously passed to the socket calls to the IPv4-mapped IPv6 address format. For example, you must convert 10.1.1.1 in the *sockaddr\_in* structure to ::ffff:10.1.1.1 in the *sockaddr\_in6* structure.

### Parameters

<i>Socket</i>	Specifies the socket descriptor (an integer) of the socket to be bound.
<i>Name</i>	Points to an address structure that specifies the address to which the socket should be bound. The <b>/usr/include/sys/socket.h</b> file defines the <b>sockaddr</b> address structure. The <b>sockaddr</b> structure contains an identifier specific to the address format and protocol provided in the <b>socket</b> subroutine.
<i>NameLength</i>	Specifies the length of the socket address structure.

## Return Values

Upon successful completion, the **bind** subroutine returns a value of 0.

If the **bind** subroutine is unsuccessful, the subroutine handler performs the following actions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see "Error Notification Object Class" in *AIX Version 6.1 Communications Programming Concepts*.

## Error Codes

The **bind** subroutine is unsuccessful if any of the following errors occurs:

Value	Description
<b>EACCES</b>	The requested address is protected, and the current user does not have permission to access it.
<b>EADDRINUSE</b>	The specified address is already in use.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The specified address is not a valid address for the address family of the specified socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EDESTADDRREQ</b>	The <i>address</i> argument is a null pointer.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the <i>UserAddress</i> space.
<b>EINVAL</b>	The socket is already bound to an address.
<b>ENOBUF</b>	Insufficient buffer space available.
<b>ENODEV</b>	The specified device does not exist.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The socket referenced by <i>Socket</i> parameter does not support address binding.

## Examples

The following program fragment illustrates the use of the **bind** subroutine to bind the name `"/tmp/zan/"` to a UNIX domain socket.

```
#include <sys/un.h>

.
.
.
struct sockaddr_un addr;
.
.
.
strcpy(addr.sun_path, "/tmp/zan/");
addr.sun_len = strlen(addr.sun_path);
addr.sun_family = AF_UNIX;
bind(s, (struct sockaddr*)&addr, SUN_LEN(&addr));
```

## Related Information

The **connect** subroutine, **getsockname** subroutine, **listen** subroutine, **socket** subroutine.

Binding Names to Sockets, Reading UNIX Datagrams Example Program, Sockets Overview, Understanding Socket Connections, and Understanding Socket Creation in *AIX Version 6.1 Communications Programming Concepts*.

---

## connect Subroutine

### Purpose

Connects two sockets.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int connect ( Socket, Name, NameLength)
int Socket;
const struct sockaddr *Name;
socklen_t NameLength;
```

### Description

The **connect** subroutine requests a connection between two sockets. The kernel sets up the communication link between the sockets; both sockets must use the same address format and protocol.

If a **connect** subroutine is issued on an unbound socket or a partially bound socket (a socket that is assigned a port number but no IP address), the system automatically binds the socket. The **connect** subroutine can be used to connect a socket to itself. This can be done, for example, by binding a socket to a local port (using **bind**) and then connecting it to the same port with a local IP address (using **connect**).

The **connect** subroutine performs a different action for each of the following two types of initiating sockets:

- If the initiating socket is **SOCK\_DGRAM**, the **connect** subroutine establishes the peer address. The peer address identifies the socket where all datagrams are sent on subsequent **send** subroutines. No connections are made by this **connect** subroutine. If the UDP socket is receiving datagrams when the **connect** subroutine is called, the subroutine will change the IP address, preventing the socket from receiving datagram packets based on the previous address.
- If the initiating socket is **SOCK\_STREAM** or **SOCK\_CONN\_DGRAM**, the **connect** subroutine attempts to make a connection to the socket specified by the *Name* parameter. Each communication space interprets the *Name* parameter differently. For **SOCK\_CONN\_DGRAM** socket type and ATM protocol, some of the ATM parameters may have been modified by the remote station, applications may query new values of ATM parameters using the appropriate socket options.
- In the case of a UNIX domain socket, a **connect** call only succeeds if the process that calls **connect** has read and write permissions on the socket file created by the **bind** call. Permissions are determined by the **umask**< value of the process that created the file.

### Implementation Specifics

### Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Name</i>	Specifies the address of target socket that will form the other end of the communication line
<i>NameLength</i>	Specifies the length of the address structure.

## Return Values

Upon successful completion, the **connect** subroutine returns a value of 0.

If the **connect** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **connect** subroutine is unsuccessful if any of the following errors occurs:

Value	Description
<b>EADDRINUSE</b>	The specified address is already in use. This error will also occur if the <b>SO_REUSEADDR</b> socket option was set and the local address (whether specified or selected by the system) is already in use.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EALREADY</b>	The socket is specified with <b>O_NONBLOCK</b> or <b>O_NDELAY</b> , and a previous connection attempt has not yet completed.
<b>EINTR</b>	The attempt to establish a connection was interrupted by delivery of a signal that was caught; the connection will be established asynchronously.
<b>EACCES</b>	Search permission is denied on a component of the path prefix or write access to the named socket is denied.
<b>ENOBUFS</b>	The system ran out of memory for an internal data structure.
<b>EOPNOTSUPP</b>	The socket referenced by <i>Socket</i> parameter does not support <b>connect</b> .
<b>EWouldBLOCK</b>	The range allocated for TCP/UDP ephemeral ports has been exhausted.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNREFUSED</b>	The attempt to connect was rejected.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINPROGRESS</b>	The socket is marked as nonblocking. The connection cannot be immediately completed. The application program can select the socket for writing during the connection process.
<b>EINVAL</b>	The specified path name contains a character with the high-order bit set.
<b>EISCONN</b>	The socket is already connected.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>ENOSPC</b>	There is no space left on a device or system table.
<b>ENOTCONN</b>	The socket could not be connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ETIMEDOUT</b>	The establishment of a connection timed out before a connection was made.
<b>EPROTOTYPE</b>	The specified address has a different type from the socket that is bound to the specified peer address.
<b>ELOOP</b>	Too many symbolic links were encountered in translating the path name in address.
<b>ENOENT</b>	A component of the path name does not name an existing file or the path name is an empty string.
<b>ENOTDIR</b>	A component of the path prefix of the path name in address is not a directory.

## Examples

The following program fragment illustrates the use of the **connect** subroutine by a client to initiate a connection to a server's socket.

```
struct sockaddr_un server;
.
.
.
connect(s,(struct sockaddr*)&server, sun_len(&server));
```

## Related Information

The **accept** subroutine, **bind** subroutine, **getsockname** subroutine, **send** subroutine, **socket**, subroutine, **socks5tcp\_connect** subroutine.

Initiating UNIXStream Connections Example Program, Sockets Overview, and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

---

## CreateloCompletionPort Subroutine

### Purpose

Creates an I/O completion port with no associated file descriptor or associates an opened socket or file with an existing or newly created I/O completion port.

### Syntax

```
#include <iocp.h>
int CreateloCompletionPort (FileDescriptor, CompletionPort, CompletionKey, ConcurrentThreads)
HANDLE FileDescriptor, CompletionPort;
DWORD CompletionKey, ConcurrentThreads;
```

### Description

The **CreateloCompletionPort** subroutine creates an I/O completion port or associates an open file descriptor with an existing or newly created I/O completion port. When creating a new I/O completion port, the *CompletionPort* parameter is set to NULL, the *FileDescriptor* parameter is set to INVALID\_HANDLE\_VALUE (-1), and the *CompletionKey* parameter is ignored.

The **CreateloCompletionPort** subroutine returns a descriptor (an integer) to the I/O completion port created or modified.

The **CreateloCompletionPort** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

**Note:** This subroutine only works with file descriptors of sockets, or regular files for use with the Asynchronous I/O (AIO) subsystem. It does not work with file descriptors of other types.

### Parameters

<i>FileDescriptor</i>	Specifies a valid file descriptor obtained from a call to the <b>socket</b> or <b>accept</b> subroutines.
<i>CompletionPort</i>	Specifies a valid I/O completion port descriptor. Specifying a <i>CompletionPort</i> parameter value of NULL causes the <b>CreateloCompletionPort</b> subroutine to create a new I/O completion port.
<i>CompletionKey</i>	Specifies an integer to serve as the identifier for completion packets generated from a particular file-completion port set.
<i>ConcurrentThreads</i>	This parameter is not implemented and is present only for compatibility purposes.

### Return Values

Upon successful completion, the **CreateloCompletionPort** subroutine returns an integer (the I/O completion port descriptor).

If the **CreateloCompletionPort** is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of NULL to the calling program.

- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable, see the link in the Related Information section of this document.

## Error Codes

The **CreateIoCompletionPort** subroutine is unsuccessful if either of the following errors occur:

<b>EBADF</b>	The I/O completion port descriptor is invalid.
<b>EINVAL</b>	The file descriptor is invalid.

## Examples

The following program fragment illustrates the use of the **CreateIoCompletionPort** subroutine to create a new I/O completion port with no associated file descriptor:

```
c = CreateIoCompletionPort (INVALID_HANDLE_VALUE, NULL, 0, 0);
```

The following program fragment illustrates the use of the **CreateIoCompletionPort** subroutine to associate file descriptor 34 (which has a newly created I/O completion port) with completion key 25:

```
c = CreateIoCompletionPort (34, NULL, 25, 0);
```

The following program fragment illustrates the use of the **CreateIoCompletionPort** subroutine to associate file descriptor 54 (which has an existing I/O completion port) with completion key 15:

```
c = CreateIoCompletionPort (54, 12, 15, 0);
```

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “ReadFile Subroutine” on page 165, “WriteFile Subroutine” on page 258, “GetQueuedCompletionStatus Subroutine” on page 101, “GetMultipleCompletionStatus Subroutine” on page 83, and “PostQueuedCompletionStatus Subroutine” on page 161.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## dn\_comp Subroutine

### Purpose

Compresses a domain name.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int dn_comp (ExpDomNam, CompDomNam, Length, DomNamPtr, LastDomNamPtr)
u_char * ExpDomNam, * CompDomNam;
int Length;
u_char ** DomNamPtr, ** LastDomNamPtr;
```

## Description

The **dn\_comp** subroutine compresses a domain name to conserve space. When compressing names, the client process must keep a record of suffixes that have appeared previously. The **dn\_comp** subroutine compresses a full domain name by comparing suffixes to a list of previously used suffixes and removing the longest possible suffix.

The **dn\_comp** subroutine compresses the domain name pointed to by the *ExpDomNam* parameter and stores it in the area pointed to by the *CompDomNam* parameter. The **dn\_comp** subroutine inserts labels into the message as the name is compressed. The **dn\_comp** subroutine also maintains a list of pointers to the message labels and updates the list of label pointers.

- If the value of the *DomNamPtr* parameter is null, the **dn\_comp** subroutine does not compress any names. The **dn\_comp** subroutine translates a domain name from ASCII to internal format without removing suffixes (compressing). Otherwise, the *DomNamPtr* parameter is the address of pointers to previously compressed suffixes.
- If the *LastDomNamPtr* parameter is null, the **dn\_comp** subroutine does not update the list of label pointers.

The **dn\_comp** subroutine is one of a set of subroutines that form the resolver. The resolver is a set of functions that perform a translation between domain names and network addresses. Global information used by the resolver subroutines resides in the **\_res** data structure. The */usr/include/resolv.h* file contains the **\_res** data structure definition.

All applications containing the **dn\_comp** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>ExpDomNam</i>	Specifies the address of an expanded domain name.
<i>CompDomNam</i>	Points to an array containing the compressed domain name.
<i>Length</i>	Specifies the size of the array pointed to by the <i>CompDomNam</i> parameter.
<i>DomNamPtr</i>	Specifies a list of pointers to previously compressed names in the current message.
<i>LastDomNamPtr</i>	Points to the end of the array specified to by the <i>CompDomNam</i> parameter.

## Return Values

Upon successful completion, the **dn\_comp** subroutine returns the size of the compressed domain name.

If unsuccessful, the **dn\_comp** subroutine returns a value of -1 to the calling program.

## Files

<i>/usr/include/resolv.h</i>	Contains global information used by the resolver subroutines.
------------------------------	---

## Related Information

The **named** daemon.

The **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_query** subroutine, **res\_search** subroutine, **res\_send** subroutine.

TCP/IP name resolution in *Networks and communication management*.

## dn\_expand Subroutine

### Purpose

Expands a compressed domain name.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int dn_expand (MessagePtr, EndOfMesOrig, CompDomNam, ExpandDomNam, Length)
u_char * MessagePtr, * EndOfMesOrig;
u_char * CompDomNam, * ExpandDomNam;
int Length;
```

### Description

The **dn\_expand** subroutine expands a compressed domain name to a full domain name, converting the expanded names to all uppercase letters. A client process compresses domain names to conserve space. Compression consists of removing the longest possible previously occurring suffixes. The **dn\_expand** subroutine restores a domain name compressed by the **dn\_comp** subroutine to its full size.

The **dn\_expand** subroutine is one of a set of subroutines that form the resolver. The resolver is a set of functions that perform a translation between domain names and network addresses. Global information used by the resolver subroutines resides in the **\_res** data structure. The **/usr/include/resolv.h** file contains the **\_res** data structure definition.

All applications containing the **dn\_expand** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>MessagePtr</i>	Specifies a pointer to the beginning of a message.
<i>EndOfMesOrig</i>	Points to the end of the original message that contains the compressed domain name.
<i>CompDomNam</i>	Specifies a pointer to a compressed domain name.
<i>ExpandDomNam</i>	Specifies a pointer to a buffer that holds the resulting expanded domain name.
<i>Length</i>	Specifies the size of the buffer pointed to by the <i>ExpandDomNam</i> parameter.

### Return Values

Upon successful completion, the **dn\_expand** subroutine returns the size of the expanded domain name.

If unsuccessful, the **dn\_expand** subroutine returns a value of -1 to the calling program.

## Files

**/etc/resolv.conf**

Defines name server and domain name constants, structures, and values.

## Related Information

The **dn\_comp** subroutine, **\_getlong** subroutine, **getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_query** subroutine, **res\_search** subroutine, **res\_send** subroutine.

TCP/IP name resolution in *Networks and communication management*.

Sockets Overview, and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## eaccept Subroutine

### Purpose

Accepts a connection on a socket to create a new socket. The **eaccept** subroutine is similar to the **accept** subroutine with the addition of the **sec\_labels\_t** structure. The **sec\_labels\_t** structure reads the Sensitivity Level (SL) that is received on the incoming connection for Trusted AIX enabled systems.

### Library

Standard C Library (**libc.a**)

### Syntax

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

int eaccept ( Socket, Address, AddressLength, Label)
int Socket;
struct sockaddr *Address;
socklen_t *AddressLength;
sec_labels_t *Label;
```

### Description

The **eaccept** subroutine extracts the first connection in the queue of pending connections, creates a new socket with the same properties as the specified socket, and allocates a new file descriptor for that socket.

If there are no connection requests in the **listen** queue, the **eaccept** subroutine performs the following actions:

- Blocks a calling socket of the **blocking** type until a connection is present.
- Returns an **EWOULDBLOCK** error code for sockets marked nonblocking.

The accepted socket cannot accept more connections, but the original socket remains open and can accept more connections.

The **eaccept** subroutine is used with only the SOCK\_STREAM socket type. If a valid *Label* parameter is specified, the SL from the incoming connection is returned to the application.

## Parameters

<i>Socket</i>	Specifies a socket created with the <b>socket</b> subroutine that is bound to an address with the <b>bind</b> or <b>ebind</b> subroutine and has issued a successful call to the <b>listen</b> subroutine.
<i>Address</i>	Specifies a result parameter that contains the address of the connecting entity as known to the communications layer. The exact format of the <i>Address</i> parameter is determined by the domain in which the communication occurs.
<i>AddressLength</i>	Specifies a parameter that initially contains the amount of space pointed to by the <i>Address</i> parameter. Upon return, the parameter contains the actual length (in bytes) of the address returned. The <b>eaccept</b> subroutine is used with <b>SOCK_STREAM</b> socket types.
<i>Label</i>	Specifies a result parameter that contains the SL received on the incoming connection.

## Return Values

Successful	a non-negative socket descriptor of the accepted socket
Unsuccessful	-1

## Error Codes

The **eaccept** subroutine is unsuccessful if one or more of the following is true:

<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EINTR</b>	The <b>eaccept</b> function was interrupted by a signal that was caught before a valid connection arrived.
<b>EINVAL</b>	The socket referenced by <i>s</i> is not currently a listen socket or has been shutdown with <b>shutdown</b> . A <b>listen</b> must be done before an <b>eaccept</b> is allowed.
<b>EMFILE</b>	The number of open file descriptors per process exceeds the system limit ( <b>OPEN_MAX</b> ).
<b>ENFILE</b>	The number of open files exceeds the allowed maximum value.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The referenced socket is not of type <b>SOCK_STREAM</b> .
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EWouldBlock</b>	The socket is marked as nonblocking, and no connections are present to be accepted.
<b>ENETDOWN</b>	The network that the socket is associated with is down.
<b>ENOTCONN</b>	The socket is not in the connected state.
<b>ECONNABORTED</b>	The client aborted the connection.
<b>EPERM</b>	The MLS MAC check failed.

---

## ebind Subroutine

### Purpose

Binds a name to a socket. Also binds a socket to the specific Sensitivity Level (SL) that is passed as a parameter.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int ebind ( Socket, Name, NameLength, Label)
int Socket;
const struct sockaddr *Name;
socklen_t NameLength;
sec_labels_t *Label;
```

## Description

The **ebind** subroutine assigns a *Name* parameter to an unnamed socket. Sockets created by the **socket** subroutine are unnamed; they are identified only by their address family. Subroutines that connect sockets either assign names or use unnamed sockets.

When a NULL pointer is passed to the *Label* parameter, then a normal multi-level port is created. However, when a valid label is passed to the *Label* parameter, a port at the specified Sensitivity Level (SL) is created. This means that only those incoming connections at the specified SL are able to connect. This also means that multiple sockets can be bound to the same port at different SLs. It is possible to create a multi-level port as well as several specific-level ports. If none of the specific SLs matches the incoming packet, then the packet port is a default multi-level port.

## Parameters

<i>Socket</i>	Specifies the socket descriptor of the socket to be bound. The socket descriptor is an integer,
<i>Name</i>	Points to an address structure that specifies the address to which the socket should be bound. The <code>/usr/include/sys/socket.h</code> file defines the <b>sockaddr</b> address structure. The <b>sockaddr</b> structure contains an identifier specific to the address format and protocol provided in the <b>socket</b> subroutine.
<i>NameLength</i>	Specifies the length of the socket address structure.
<i>Label</i>	Specifies the Sensitivity Label associated with the socket.

## Return Values

Successful	0
Unsuccessful	-1

## Error Codes

The **ebind** subroutine is unsuccessful if any of the following errors occurs:

Value	Description
<b>EACCES</b>	The requested address is protected, and the current user does not have permission to access it.
<b>EADDRINUSE</b>	The specified address is already in use.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The specified address is not a valid address for the address family of the specified socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EDESTADDRREQ</b>	The <i>address</i> argument is a null pointer.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINVAL</b>	The socket is already bound to an address.
<b>ENOBUF</b>	Insufficient buffer space available.
<b>ENODEV</b>	The specified device does not exist.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The socket referenced by the <i>Socket</i> parameter does not support address binding.

---

## econnect Subroutine

### Purpose

Connects two sockets. The **econnect** subroutine is similar to the **connect** subroutine with the addition of the **sec\_labels\_t** pointer. The **sec\_labels\_t** pointer indicates the Sensitivity Level (SL) of the outgoing connection request.

### Library

Standard C Library (**libc.a**)

### Syntax

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

int econnect ( Socket, Name, NameLength, Label)
int Socket;
const struct sockaddr *Name;
socklen_t NameLength;
sec_labels_t *Label;
```

### Description

The **econnect** subroutine requests a connection between two sockets, similar to the **connect** subroutine. The kernel sets up the communication link between the sockets; both sockets must use the same address format and protocol.

The SL specified by the *Label* parameter is the SL of the outgoing request. The requested SL must be dominated by the current clearance or must have appropriate privileges to clear the MAC check.

### Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Name</i>	Specifies the address of the target socket that will form the other end of the communication line.
<i>NameLength</i>	Specifies the length of the address structure.
<i>Label</i>	Specifies the SL of the outgoing connection request.

### Return Values

Successful	0
Unsuccessful	-1

### Error Codes

The **econnect** subroutine is unsuccessful if any of the following errors occurs:

Value	Description
<b>EADDRINUSE</b>	The specified address is already in use.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EALREADY</b>	The socket is specified with <b>O_NONBLOCK</b> or <b>O_NDELAY</b> , and a previous connection attempt has not yet completed.
<b>EINTR</b>	The attempt to establish a connection was interrupted by delivery of a signal that was caught; the connection will be established asynchronously.
<b>EACCES</b>	Search permission was denied on a component of the path prefix or write access to the named socket was denied.

Value	Description
<b>ENOBUFS</b>	The system has run out of memory for an internal data structure.
<b>EOPNOTSUPP</b>	The socket referenced by the <i>Socket</i> parameter does not support the <b>connect</b> subroutine.
<b>EWOULDBLOCK</b>	The range allocated for TCP/UDP ephemeral ports has been exhausted.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNREFUSED</b>	The attempt to connect was rejected.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINPROGRESS</b>	The socket is marked as nonblocking. The connection cannot be immediately completed. The application program can select the socket for writing during the connection process.
<b>EINVAL</b>	The specified path name contains a character with the high-order bit set.
<b>EISCONN</b>	The socket is already connected.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>ENOSPC</b>	There is no space left on a device or system table.
<b>ENOTCONN</b>	The socket could not be connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ETIMEDOUT</b>	The establishment of a connection times out before a connection is made.
<b>EPERM</b>	The Trusted AIX MAC check failed.

---

## endhostent Subroutine

### Purpose

Closes the **/etc/hosts** file.

### Library

Standard C Library (**libc.a**)  
**(libbind)**  
**(libnis)**  
**(liblocal)**

### Syntax

```
#include <netdb.h>
endhostent ()
```

### Description

When using the **endhostent** subroutine in DNS/BIND name service resolution, **endhostent** closes the TCP connection which the **sethostent** subroutine set up.

When using the **endhostent** subroutine in NIS name resolution or to search the **/etc/hosts** file, **endhostent** closes the **/etc/hosts** file.

**Note:** If a previous **sethostent** subroutine is performed and the *StayOpen* parameter does not equal 0, the **endhostent** subroutine closes the **/etc/hosts** file. Run a second **sethostent** subroutine with the *StayOpen* value equal to 0 in order for a following **endhostent** subroutine to succeed. Otherwise, the **/etc/hosts** file closes on an **exit** subroutine call .

### Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name service ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

The **gethostbyaddr** subroutine, **gethostbyname** subroutine, **sethostent** subroutine **gethostent** subroutine.

Sockets Overview and Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## endhostent\_r Subroutine

### Purpose

Closes the **/etc/hosts** file.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include <netdb.h>
```

```
void endhostent_r (struct hostent_data *ht_data);
```

### Description

When using the **endhostent\_r** subroutine in DNS/BIND name service resolution, **endhostent\_r** closes the TCP connection which the **sethostent\_r** subroutine set up.

When using the **endhostent\_r** subroutine in NIS name resolution or to search the **/etc/hosts** file, **endhostent\_r** closes the **/etc/hosts** file.

**Note:** If a previous **sethostent\_r** subroutine is performed and the *StayOpen* parameter does not equal 0, then the **endhostent\_r** subroutine closes the **/etc/hosts** file. Run a second **sethostent\_r** subroutine with the *StayOpen* value equal to 0 in order for a following **endhostent\_r** subroutine to succeed. Otherwise, the **/etc/hosts** file closes on an **exit** subroutine call .

### Parameters

<i>ht_data</i>	Points to the <b>hostent_data</b> structure
----------------	---

### Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name service ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

“gethostbyaddr\_r Subroutine” on page 75, “gethostbyname\_r Subroutine” on page 78, “sethostent\_r Subroutine” on page 207, and “gethostent\_r Subroutine” on page 80.

---

## endnetent Subroutine

### Purpose

Closes the **/etc/networks** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void endnetent ( )
```

### Description

The **endnetent** subroutine closes the **/etc/networks** file. Calls made to the **getnetent**, **getnetbyaddr**, or **getnetbyname** subroutine open the **/etc/networks** file.

All applications containing the **endnetent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Return Values

If a previous **setnetent** subroutine has been performed and the *StayOpen* parameter does not equal 0, then the **endnetent** subroutine will not close the **/etc/networks** file. Also, the **setnetent** subroutine does not indicate that it closed the file. A second **setnetent** subroutine has to be issued with the *StayOpen* parameter equal to 0 in order for a following **endnetent** subroutine to succeed. If this is not done, the **/etc/networks** file must be closed with the **exit** subroutine.

### Examples

To close the **/etc/networks** file, type:

```
endnetent();
```

### Files

<b>/etc/networks</b>	Contains official network names.
----------------------	----------------------------------

### Related Information

The **exit** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **setnetent** subroutine.

Sockets Overview, Understanding Network Address Translation, and List of Socket Programming References in *AIX Version 6.1 Communications Programming Concepts*.

---

## endnetent\_r Subroutine

### Purpose

Closes the **/etc/networks** file.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <netdb.h>
```

```
void endnetent_r (net_data)
struct netent_data *net_data;
```

## Description

The **endnetent\_r** subroutine closes the **/etc/networks** file. Calls made to the **getnetent\_r**, **getnetbyaddr\_r**, or **getnetbyname\_r** subroutine open the **/etc/networks** file.

## Parameters

<i>net_data</i>	Points to the <b>netent_data</b> structure.
-----------------	---

## Files

<b>/etc/networks</b>	Contains official network names.
----------------------	----------------------------------

## Related Information

“getnetbyaddr\_r Subroutine” on page 88, “getnetbyname\_r Subroutine” on page 90, “getnetent\_r Subroutine” on page 92, and “setnetent\_r Subroutine” on page 211.

---

## endnetgrent\_r Subroutine

### Purpose

Handles the group network entries.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <netdb.h>
void endnetgrent_r (void **ptr)
```

## Description

The **setnetgrent\_r** subroutine establishes the network group from which the **getnetgrent\_r** subroutine will obtain members, and also restarts calls to the **getnetgrent\_r** subroutine from the beginning of the list. If the previous **setnetgrent\_r** call was to a different network group, an **endnetgrent\_r** call is implied.

The **endnetgrent\_r** subroutine frees the space allocated during the **getnetgrent\_r** calls.

## Parameters

<i>ptr</i>	Keeps the function threadsafe.
------------	--------------------------------

## Files

<b>/etc/netgroup</b>	Contains network groups recognized by the system.
<b>/usr/include/netdb.h</b>	Contains the network database structures.

## Related Information

“getnetgrent\_r Subroutine” on page 93, and “setnetgrent\_r Subroutine” on page 211.

---

## endprotoent Subroutine

### Purpose

Closes the **/etc/protocols** file.

### Library

Standard C Library (**libc.a**)

### Syntax

**void endprotoent (void)**

### Description

The **endprotoent** subroutine closes the **/etc/protocols** file.

Calls made to the **getprotoent** subroutine, **getprotobyname** subroutine, or **getprotobynumber** subroutine open the **/etc/protocols** file. An application program can use the **endprotoent** subroutine to close the **/etc/protocols** file.

All applications containing the **endprotoent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Return Values

If a previous **setprotoent** subroutine has been performed and the *StayOpen* parameter does not equal 0, the **endprotoent** subroutine will not close the **/etc/protocols** file. Also, the **setprotoent** subroutine does not indicate that it closed the file. A second **setprotoent** subroutine has to be issued with the *StayOpen* parameter equal to 0 in order for a following **endprotoent** subroutine to succeed. If this is not done, the **/etc/protocols** file closes on an **exit** subroutine.

### Examples

To close the **/etc/protocols** file, type:

```
endprotoent();
```

### Files

<b>/etc/protocols</b>	Contains protocol names.
-----------------------	--------------------------

## Related Information

The **exit** subroutine, **getprotobynumber** subroutine, **getprotobyname** subroutine, **getprotoent** subroutine, **setprotoent** subroutine.

Sockets Overview, and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## endprotoent\_r Subroutine

### Purpose

Closes the **/etc/protocols** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
void endprotoent_r(proto_data);
struct protoent_data *proto_data;
```

### Description

The **endprotoent\_r** subroutine closes the **/etc/protocols** file, which is opened by the calls made to the **getprotoent\_r** subroutine, **getprotobyname\_r** subroutine, or **getprotobyname\_r** subroutine.

### Parameters

<i>proto_data</i>	Points to the <b>protoent_data</b> structure
-------------------	--

### Files

<b>/etc/protocols</b>	Contains protocol names.
-----------------------	--------------------------

### Related Information

“getprotobyname\_r Subroutine” on page 98, “getprotobyname\_r Subroutine” on page 96, “getprotoent\_r Subroutine” on page 100, and “setprotoent\_r Subroutine” on page 213.

---

## endservent Subroutine

### Purpose

Closes the **/etc/services** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void endservent ( )
```

### Description

The **endservent** subroutine closes the **/etc/services** file. A call made to the **getservent** subroutine, **getservbyname** subroutine, or **getservbyport** subroutine opens the **/etc/services** file. An application program can use the **endservent** subroutine to close the **/etc/services** file.

All applications containing the **endservent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the **BSD libbsd.a** library.

## Return Values

If a previous **setservent** subroutine has been performed and the *StayOpen* parameter does not equal 0, then the **endservent** subroutine will not close the **/etc/services** file. Also, the **setservent** subroutine does not indicate that it closed the file. A second **setservent** subroutine has to be issued with the *StayOpen* parameter equal to 0 in order for a following **endservent** subroutine to succeed. If this is not done, the **/etc/services** file closes on an **exit** subroutine.

## Examples

To close the **/etc/services** file, type:

```
endservent ();
```

## Files

**/etc/services**

Contains service names.

## Related Information

The **endprotoent** subroutine, **exit** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **getservent** subroutine, **setprotoent** subroutine, **setservent** subroutine.

Sockets Overview, Understanding Network Address Translation, and List of Socket Programming References in *AIX Version 6.1 Communications Programming Concepts*.

---

## endservent\_r Subroutine

### Purpose

Closes the **/etc/services** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void endservent_r(serv_data)
struct servent_data *serv_data;
```

### Description

The **endservent\_r** subroutine closes the **/etc/services** file, which is opened by a call made to the **getservent\_r** subroutine, **getservbyname\_r** subroutine, or **getservbyport\_r** subroutine opens the **/etc/services** file.

### Parameters

*serv\_data*

Points to the **servent\_data** structure

## Examples

To close the **/etc/services** file, type:

```
endservent_r(serv_data);
```

## Files

**/etc/services**

Contains service names.

## Related Information

“setserverent\_r Subroutine” on page 215, “getserverent\_r Subroutine” on page 109, “getservbyport Subroutine” on page 105, and “getservbyname\_r Subroutine” on page 104.

---

## erecv, erecvmsg, erecvfrom, enrecvmsg, or enrecvfrom Subroutine

### Purpose

Allows applications to receive messages from sockets along with the Sensitivity Level (SL).

### Library

Standard C Library (**libc.a**)

Trusted AIX Sensitivity Label Library (**libmls.a**)

### Syntax

```
#include <sys/socket.h>
#include <sys/mac.h>
int erecv (Socket, Buffer, Length, Flags, Label)
int Socket;
void * Buffer;
size_t Length;
int Flags;
sec_labels_t *Label;

int erecvmsg ( Socket, Message, Flags, Label)
int Socket;
struct msghdr Message [ ];
int Flags;
sec_labels_t *Label;

ssize_t erecvfrom (Socket, Buffer, Length, Flags, From, FromLength, Label)
int Socket;
void * Buffer;
size_t Length;
int Flags;
struct sockaddr * From;
socklen_t * FromLength;
sec_labels_t *Label;

int enrecvmsg (Socket, Message, Flags, Label)
int Socket;
struct msghdr Message [ ];
int Flags;
sec_labels_t *Label;

ssize_t enrecvfrom (Socket, Buffer, Length, Flags, From, FromLength, Label)
int Socket;
void *Buffer;
size_t Length;
int Flags;
struct sockaddr *From;
socklen_t *FromLength;
sec_labels_t *Label;
```

## Description

The **erecv**, **erecvmsg**, **erecvfrom**, **enrecvmsg**, and **enrecvfrom** subroutines work exactly like the **recv**, **recvmsg**, **recvfrom**, **nrecvmsg**, and **nrecvfrom** subroutines respectively, except that the **erecv**, **erecvmsg**, **erecvfrom**, **enrecvmsg**, and **enrecvfrom** subroutines allow the application to retrieve the SL from the received data by providing a valid *Label* parameter.

If no messages are available at the socket, the **erecv**, **erecvmsg**, **erecvfrom**, **enrecvmsg**, and **enrecvfrom** subroutines wait for a message to arrive, unless the socket is nonblocking. If a socket is nonblocking, the system returns an error.

## Parameters

<i>Socket</i>	Specifies the socket descriptor.
<i>Buffer</i>	Specifies the address where the message is placed.
<i>Length</i>	Specifies the size of the <i>Buffer</i> parameter.
<i>Flags</i>	Points to a value controlling the message reception. The <b>/usr/include/sys/socket.h</b> file defines the <i>Flags</i> parameter. The argument to receive a call is formed by the logical OR operation with one or more of the following values:  <b>MSG_OOB</b> Processes out-of-band data. The significance of out-of-band data is protocol dependent.  <b>MSG_PEEK</b> Peeks at incoming data. The data continues to be treated as unread and will be read by the next call to the <b>erecv</b> , <b>erecvmsg</b> , <b>erecvfrom</b> , <b>enrecvmsg</b> , or <b>enrecvfrom</b> subroutine or a similar subroutine.  <b>MSG_WAITALL</b> Requests that the subroutine does not return until the requested number of bytes are read. The subroutine can return fewer bytes than the requested number if a signal is caught, the connection is terminated, or an error is pending for the socket. The subroutine can also return fewer bytes when the SL information across the data stream is different. Only those bytes that have the same SL information are returned to the user.
<i>Message</i>	Points to the address of the <b>msghdr</b> structure, which contains both the address for the incoming message and the space for the sender address.
<i>From</i>	Points to a socket structure, containing the address of the source.
<i>FromLength</i>	Specifies the length of the address of the sender or of the source.
<i>Label</i>	Specifies a result parameter that contains the SL from the received data.

## Return Values

Upon successful completion, the subroutines return the length of the message in bytes.

When an error occurs, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Returns a value of 0 if the connection disconnects (in case of connected sockets).
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **erecv**, **erecvmsg**, **erecvfrom**, **enrecvmsg**, or **enrecvfrom** subroutine is unsuccessful if any of the following errors occurs:

<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	The remote peer forced the connection to be closed.

<b>EFAULT</b>	The data was directed into a nonexistent or protected part of the process address space. (The <i>Buffer</i> parameter is not valid.)
<b>EINTR</b>	A signal interrupted the <b>erecv</b> , <b>erecvmsg</b> , <b>erecvfrom</b> , <b>enrecvmsg</b> , or <b>enrecvfrom</b> subroutine before any data is available.
<b>EINVAL</b>	The <b>MSG_OOB</b> value was set and no out-of-band data was available.
<b>ENOBUF</b>	Insufficient resources are available in the system to perform the operation.
<b>ENOTCONN</b>	A receiving operation was attempted on a <b>SOCK_STREAM</b> socket that was not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The <b>MSG_OOB</b> value is set for a <b>SOCK_DGRAM</b> socket or any <b>AF_UNIX</b> socket.
<b>ETIMEDOUT</b>	The connection timed out during connection establishment, or there was a transmission timeout on an active connection.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.
<b>EACCES</b>	The MLS MAC check failed.

---

## esend, esendto, or esendmsg Subroutine

### Purpose

Allows an application to send messages on a socket with the Sensitivity Level (SL) different from that of its own.

### Library

Standard C Library (**libc.a**)

Trusted AIX Sensitivity Label Library (**libmls.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/socketvar.h>
#include <sys/mac.h>
#include <sys/socket.h>
int esend (Socket, Message, Length, Flags, Label)
int Socket;
const void * Message;
size_t Length;
int Flags;
sec_labels_t *Label;

int esendmsg ( Socket, Message, Flags, Label)
int Socket;
const struct msghdr Message [ ];
int Flags;
sec_labels_t *Label;

int esendto (Socket, Message, Length, Flags, To, ToLength, Label)
int Socket;
const void * Message;
size_t Length;
int Flags;
const struct sockaddr * To;
socklen_t ToLength;
sec_labels_t *Label;
```

### Description

The **esend**, **esendmsg**, and **esendto** subroutines work exactly like **send**, **sendmsg** and **sendto** subroutines respectively, except that the **esend**, **esendmsg**, and **esendto** subroutines allow applications to associate a Sensitivity Level different from their own to the outgoing data through the *Label* parameter.

The **esend** subroutine can be used on connected sockets only. The **esendto** and **esendmsg** subroutines can be used with connected or unconnected sockets.

For **SOCK\_STREAM** socket types, when the SL is changed between subsequent send operations, the application is blocked until the pending data on the socket buffer can be flushed. If the socket is marked as nonblocking type and there is pending data on the socket buffer, an error is returned.

## Parameters

<i>Socket</i>	Specifies a unique name for the socket.
<i>Message</i>	Points to the address of the message or the <b>msghdr</b> structure containing the message to send.
<i>Length</i>	Specifies the length of the message in bytes.
<i>Flags</i>	Allows the sender to control the transmission of the message.
	<b>MSG_OOB</b> Processes out-of-band data on sockets that support <b>SOCK_STREAM</b> communication.
	<b>MSG_DONTROUTE</b> Sends without using routing tables.
	<b>MSG_MPEG2</b> Indicates that this block is a MPEG2 block. This value is valid <b>SOCK_CONN_DGRAM</b> socket types only.
<i>To</i>	Specifies the destination address for the message. The destination address is a <b>sockaddr</b> structure defined in the <b>/usr/include/sys/socket.h</b> file.
<i>ToLength</i>	Specifies the size of the destination address.
<i>Label</i>	Specifies the SL to be used on the outgoing data.

## Return Values

Upon successful completion, the **esend**, **esendmsg**, or **esendto** subroutine returns the number of characters sent.

If errors occur, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **esend**, **esendmsg**, or **esendto** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EACCES</b>	Write access to the named socket is denied, or the socket trying to send a broadcast packet does not have broadcast capability, or the MLS MAC check failed.
<b>EADDRNOTAVAIL</b>	The specified address is not valid.
<b>EAFNOSUPPORT</b>	The specified address is not a valid address for the address family of this socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	A connection was forcibly closed by a peer.
<b>EDESTADDRREQ</b>	The socket is not in connection mode and no peer address is set.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EHOSTUNREACH</b>	The destination host cannot be reached.
<b>EINTR</b>	A signal interrupted the <b>esend</b> , <b>esendmsg</b> , or <b>esendto</b> subroutine before any data was transmitted.
<b>EINVAL</b>	The <i>Length</i> parameter is not valid.
<b>EISCONN</b>	A <b>SOCK_DGRAM</b> socket is already connected.
<b>EMSGSIZE</b>	The message is too large to be sent all at once, as the socket requires.

Error	Description
ENETUNREACH	The destination network is not reachable.
ENOBUFS	Insufficient resources were available in the system to perform the operation.
ENOENT	The path name does not contain an existing file, or the path name is an empty string.
ENOMEM	The available data space in memory is not large enough to hold group or ACL information.
ENOTSOCK	The <i>Socket</i> parameter refers to a file, not a socket.
EOPNOTSUPP	The <i>Socket</i> parameter is associated with a socket that does not support one or more of the values set in the <i>Flags</i> parameter.
EPIPE	An attempt was made to send on a socket that was connected, but the connection was shut down either by the remote peer or by this side of the connection. If the socket is of type <b>SOCK_STREAM</b> , the <b>SIGPIPE</b> signal is generated for the calling process.
EWOULDBLOCK	The socket is marked nonblocking, and no connections are present to be accepted. Or a sending operation was attempted with different SLs while there was pending data on the socket buffer, and the socket was marked nonblocking

---

## ether\_ntoa, ether\_aton, ether\_ntohost, ether\_hostton, or ether\_line Subroutine

### Purpose

Maps 48-bit Ethernet numbers.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/socket.h>
#include <net/if.h>
#include <netinet/in.h>
#include <netinet/if_ether.h>
```

```
char *ether_ntoa (EthernetNumber)
struct ether_addr * EthernetNumber;
```

```
struct ether_addr *ether_aton( String);
char *string
```

```
int *ether_ntohost (HostName, EthernetNumber)
char * HostName;
struct ether_addr *EthernetNumber;
int *ether_hostton (HostName, EthernetNumber)
char *HostName;
struct ether_addr *EthernetNumber;
```

```
int *ether_line (Line, EthernetNumber, HostName)
char * Line, *HostName;
struct ether_addr *EthernetNumber;
```

### Description

**Attention:** Do not use the **ether\_ntoa** or **ether\_aton** subroutine in a multithreaded environment.

The **ether\_ntoa** subroutine maps a 48-bit Ethernet number pointed to by the *EthernetNumber* parameter to its standard ASCII representation. The subroutine returns a pointer to the ASCII string. The

representation is in the form *x:x:x:x:x*: where *x* is a hexadecimal number between 0 and ff. The **ether\_aton** subroutine converts the ASCII string pointed to by the *String* parameter to a 48-bit Ethernet number. This subroutine returns a null value if the string cannot be scanned correctly.

The **ether\_ntohost** subroutine maps a 48-bit Ethernet number pointed to by the *EthernetNumber* parameter to its associated host name. The string pointed to by the *HostName* parameter must be long enough to hold the host name and a null character. The **ether\_hostton** subroutine maps the host name string pointed to by the *HostName* parameter to its corresponding 48-bit Ethernet number. This subroutine modifies the Ethernet number pointed to by the *EthernetNumber* parameter.

The **ether\_line** subroutine scans the line pointed to by *line* and sets the hostname pointed to by the *HostName* parameter and the Ethernet number pointed to by the *EthernetNumber* parameter to the information parsed from *LINE*.

## Parameters

<i>EthernetNumber</i>	Points to an Ethernet number.
<i>String</i>	Points to an ASCII string.
<i>HostName</i>	Points to a host name.
<i>Line</i>	Points to a line.

## Return Values

0	Indicates that the subroutine was successful.
non-zero	Indicates that the subroutine was not successful.

## Files

<i>/etc/ethers</i>	Contains information about the known (48-bit) Ethernet addresses of hosts on the Internet.
--------------------	--

## Related Information

Subroutines Overview and List of Multithread Subroutines in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## FrcaCacheCreate Subroutine

### Purpose

Creates a cache instance within the scope of a Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCacheCreate ( CacheHandle, FrcaHandle, CacheSpec);
int32_t * CacheHandle;
int32_t FrcaHandle;
frca_cache_create_t * CacheSpec;
```

## Description

The **FrcaCacheCreate** subroutine creates a cache instance for an FRCA instance that has already been configured. Multiple caches can be created for an FRCA instance. Cache handles are unique only within the scope of the FRCA instance.

## Parameters

<i>CacheHandle</i>	Returns a handle that is required by the other cache-related subroutines of the FRCA API to refer to the newly created FRCA cache instance.
<i>FrcaHandle</i>	Identifies the FRCA instance for which the cache is created.
<i>CacheSpec</i>	Points to a <b>frca_ctrl_create_t</b> structure, which specifies the characteristics of the cache to be created. The structure contains the following members:

```
uint32_t  cacheType;  
uint32_t  nMaxEntries;
```

**Note:** Structure members do not necessarily appear in this order.

*cacheType*

Specifies the type of the cache instance. This field must be set to **FCTRL\_SERVERTYPE\_HTTP**.

*nMaxEntries*

Specifies the maximum number of entries allowed for the cache instance.

## Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

<b>EINVAL</b>	The <i>CacheHandle</i> or the <i>CacheSpec</i> parameter is zero or the <i>CacheSpec</i> parameter is not of the correct type <b>FCTRL_CACHETYPE_HTTP</b> .
<b>EFAULT</b>	The <i>CacheHandle</i> or the <i>CacheSpec</i> point to an invalid address.
<b>ENOENT</b>	The <i>FrcaHandle</i> parameter is invalid.

## Related Information

The **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCacheDelete Subroutine

### Purpose

Deletes a cache instance within the scope of a Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

## Syntax

```
#include <frca.h>
int32_t FrcaCacheDelete ( CacheHandle, FrcaHandle);
int32_t CacheHandle;
int32_t FrcaHandle;
```

## Description

The **FrcaCacheDelete** subroutine deletes a cache instance and releases any associated resources.

## Parameters

<i>CacheHandle</i>	Identifies the cache instance that is to be deleted.
<i>FrcaHandle</i>	Identifies the FRCA instance to which the cache instance belongs.

## Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

**ENOENT** The *CacheHandle* or the *FrcaHandle* parameter is invalid.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCacheLoadFile Subroutine

### Purpose

Loads a file into a cache associated with a Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCacheLoadFile ( CacheHandle, FrcaHandle, FileSpec, AssocData);
int32_t CacheHandle;
int32_t FrcaHandle;
frca_filespec_t * FileSpec;
frca_assocdata_t * AssocData;
```

### Description

The **FrcaCacheLoadFile** subroutine loads a file into an existing cache instance for an previously configured FRCA instance.

## Parameters

*CacheHandle* Identifies the cache instance to which the new entry should be added.  
*FrcaHandle* Identifies the FRCA instance to which the cache instance belongs.  
*FileSpec* Points to a **frca\_loadfile\_t** structure, which specifies characteristics used to identify the cache entry that is to be loaded into the given cache. The structure contains the following members:

```
uint32_t  cacheEntryType;  
char * fileName;  
char * virtualHost;  
char * searchKey;
```

**Note:** Structure members do not necessarily appear in this order.

*cacheEntryType*

Specifies the type of the cache entry. This field must be set to FCTRL\_CET\_HTTPFILE.

*fileName*

Specifies the absolute path to the file that is providing the contents for the new cache entry.

*virtualHost*

Specifies a virtual host name that is being served by the FRCA instance.

*searchKey*

Specifies the key that the cache entry can be found under by the FRCA instance when it processes an intercepted request. For the HTTP GET engine, the search key is identical to the *abs\_path* part of the HTTP URL according to section 3.2.2 of RFC 2616. For example, the search key corresponding to the URL `http://www.mydomain/welcome.html` is `/welcome.html`.

**Note:** If a cache entry with the same type, file name, virtual host, and search key already exists and the file has not been modified since the existing entry was created, the load request succeeds without any effect. If the entry exists and the file's contents have been modified since being loaded into the cache, the cache entry is updated. If the entry exists and the file's contents have not changed, but any of the settings of the HTTP header fields change, the existing entry must be unloaded first.

*AssocData* Points to a **frca\_assocdata\_t** structure, which specifies additional information to be associated with the contents of the given cache entry. The structure contains the following members:

```
uint32_t assocDataType;
char * cacheControl;
char * contentType;
char * contentEncoding;
char * contentLanguage;
char * contentCharset;
```

**Note:** Structure members do not necessarily appear in this order.

*assocDataType*

Specifies the type of data that is associated with the given cache entry.

*cacheControl*

Specifies the settings of the corresponding HTTP header field according to RFC 2616.

*contentType*

Specifies the settings of the corresponding HTTP header field according to RFC 2616.

*contentEncoding*

Specifies the settings of the corresponding HTTP header field according to RFC 2616.

*contentLanguage*

Specifies the settings of the corresponding HTTP header field according to RFC 2616.

*contentCharset*

Specifies the settings of the corresponding HTTP header field according to RFC 2616.

## Return Values

0

The subroutine completed successfully.

-1

The subroutine failed. The global variable *errno* is set to indicate the specific type of error.

## Error Codes

**EINVAL**

The *FileSpec* or the *AssocData* parameter is zero or are not of the correct type or any of the *fileName* or the *searchKey* components are zero or the size of the file is zero.

**EFAULT**

The *FileSpec* or the *AssocData* parameter or one of their components points to an invalid address.

**ENOMEM**

The FRCA or NBC subsystem is out of memory.

**EFBIG**

The content of the cache entry failed to load into the NBC. Check network options **nbc\_limit**, **nbc\_min\_cache**, and **nbc\_max\_cache**.

**ENOTREADY**

The kernel extension is currently being loaded or unloaded.

**ENOENT**

The *CacheHandle* or the *FrcaHandle* parameter is invalid.

## Related Information

The **FrcaCacheCreate** subroutine, **FrcaCacheDelete** subroutine, **FrcaCacheUnloadFile** subroutine, **FrcaCtrlCreate** subroutine, **FrcaCtrlDelete** subroutine, **FrcaCtrlLog** subroutine, **FrcaCtrlStart** subroutine, **FrcaCtrlStop** subroutine.

---

## FrcaCacheUnoadFile Subroutine

### Purpose

Removes a cache entry from a cache that is associated with a Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCacheUnoadFile ( CacheHandle, FrcaHandle, FileSpec);
int32_t CacheHandle;
int32_t FrcaHandle;
frca_filespec_t * FileSpec;
```

### Description

The **FrcaCacheUnoadFile** subroutine removes a cache entry from an existing cache instance for an previously configured FRCA instance.

### Parameters

<i>CacheHandle</i>	Identifies the cache instance from which the entry should be removed.
<i>FrcaHandle</i>	Identifies the FRCA instance to which the cache instance belongs.
<i>FileSpec</i>	Points to a <b>frca_loadfile_t</b> structure, which specifies characteristics used to identify the cache entry that is to be removed from the given cache. The structure contains the following members:

```
uint32_t cacheEntryType;
char * fileName;
char * virtualHost;
char * searchKey;
```

**Note:** Structure members do not necessarily appear in this order.

*cacheEntryType*

Specifies the type of the cache entry. This field must be set to FCTRL\_CET\_HTTPFILE.

*fileName*

Specifies the absolute path to the file that is to be removed from the cache.

*virtualHost*

Specifies a virtual host name that is being served by the FRCA instance.

*searchKey*

Specifies the key under which the cache entry can be found.

**Note:** The **FrcaCacheUnoadFile** subroutine succeeds if a cache entry with the same type, file name, virtual host, and search key does not exist. This subroutine fails if the file associated with *fileName* does not exist or if the calling process does not have sufficient access permissions.

### Return Values

0	The subroutine completed successfully.
---	--

-1

The subroutine failed. The global variable *errno* is set to indicate the specific type of error.

## Error Codes

**EINVAL**

The *FileSpec* parameter is zero or the *cacheEntryType* component is not set to **FCTRL\_CET\_HTTPFILE** or the *searchKey* component is zero or the *fileName* is '/' or the *fileName* is not an absolute path.

**EFAULT**

The *FileSpec* parameter or one of the components points to an invalid address.

**EACCES**

Access permission is denied on the *fileName*.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCtrlCreate Subroutine

### Purpose

Creates a Fast Response Cache Accelerator (FRCA) control instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCtrlCreate ( FrcaHandle, InstanceSpec);
int32_t * FrcaHandle;
frca_ctrl_create_t * InstanceSpec;
```

### Description

The **FrcaCtrlCreate** subroutine creates and configures an FRCA instance that is associated with a previously configured TCP listen socket. TCP connections derived from the TCP listen socket are intercepted by the FRCA instance and, if applicable, adequate responses are generated by the in-kernel code on behalf of the user-level application.

The only FRCA instance type that is currently supported handles static GET requests as part of the Hypertext Transfer Protocol (HTTP).

### Parameters

*FrcaHandle* Returns a handle that is required by the other FRCA API subroutines to refer to the newly configured FRCA instance.

*InstanceSpec* Points to a **frca\_ctrl\_create\_t** structure, which specifies the parameters used to configure the newly created FRCA instance. The structure contains the following members:

```
uint32_t  serverType;  
char *    serverName;  
uint32_t  nListenSockets;  
uint32_t * ListenSockets;  
uint32_t  flags;  
uint32_t  nMaxConnections;  
uint32_t  nLogBufs;  
char *    logFile;
```

**Note:** Structure members do not necessarily appear in this order.

*serverType*

Specifies the type for the FRCA instance. This field must be set to **FCTRL\_SERVERTYPE\_HTTP**.

*serverName*

Specifies the value to which the HTTP header field is set.

*nListenSocket*

Specifies the number of listen socket descriptors pointed to by *listenSockets*.

*listenSocket*

Specifies the TCP listen socket that the FRCA instance should be configured to intercept.

**Note:** The TCP listen socket must exist and the **SO\_KERNACCEPT** socket option must be set at the time of calling the **FrcaCtrlCreate** subroutine.

*flags*

Specifies the logging format, the initial state of the logging subsystem, and whether responses generated by the FRCA instance should include the **Server:** HTTP header field. The valid flags are as follows:

**FCTRL\_KEEPLIVE**

**FCTRL\_LOGFORMAT**

**FCTRL\_LOGFORMAT\_ECLF**

**FCTRL\_LOGFORMAT\_VHOST**

**FCTRL\_LOGMODE**

**FCTRL\_LOGMODE\_ON**

**FCTRL\_SENDSERVERHEADER**

*nMaxConnections*

Specifies the maximum number of intercepted connections that are allowed at any given point in time.

*nLogBufs*

Specifies the number of preallocated logging buffers used for logging information about HTTP GET requests that have been served successfully.

*logFile*

Specifies the absolute path to a file used for appending logging information. The HTTP GET engine uses *logFile* as a base name and appends a sequence number to it to generate the actual file name. Whenever the size of the current log file exceeds the threshold of approximately 1 gigabyte, the sequence number is incremented by 1 and the logging subsystem starts appending to the new log file.

**Note:** The FRCA instance creates the log file, but not the path to it. If the path does not exist or is not accessible, the FRCA instance reverts to the default log file **/tmp/frca.log**.

## Return Values

0

The subroutine completed successfully.

-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.
----	--

## Error Codes

EINVAL	The <i>FrcaHandle</i> or the <i>InstanceSpec</i> parameter is zero or is not of the correct type or the <i>listenSockets</i> components do not specify any socket descriptors.
EFAULT	The <i>FrcaHandle</i> or the <i>InstanceSpec</i> or a component of the <i>InstanceSpec</i> points to an invalid address.
ENOTREADY	The kernel extension is currently being loaded or unloaded.
ENOTSOCK	A TCP listen socket does not exist.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCtrlDelete Subroutine

### Purpose

Deletes a Fast Response Cache Accelerator (FRCA) control instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCtrlDelete ( FrcaHandle);
int32_t * FrcaHandle;
```

### Description

The **FrcaCtrlDelete** subroutine deletes an FRCA instance and releases any associated resources.

The only FRCA instance type that is currently supported handles static GET requests as part of the Hypertext Transfer Protocol (HTTP).

### Parameters

<i>FrcaHandle</i>	Identifies the FRCA instance on which this operation is performed.
-------------------	--

### Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

<b>ENOENT</b>	The <i>FrcaHandle</i> parameter is invalid.
<b>ENOTREADY</b>	The FRCA control instance is in an undefined state.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCtrlLog Subroutine

### Purpose

Modifies the behavior of the logging subsystem.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCtrlLog ( FrcaHandle,  Flags);
int32_t FrcaHandle;
uint32_t Flags;
```

### Description

The **FrcaCtrlLog** subroutine modifies the behavior of the logging subsystem for the Fast Response Cache Accelerator (FRCA) instance specified. Modifiable attributes are the logging mode, which can be turned on or off, and the logging format, which defaults to the HTTP Common Log Format (CLF). The logging format can be changed to Extended Common Log Format (ECLF) and can be set to include virtual host information.

The only FRCA instance type that is currently supported handles static GET requests as part of the Hypertext Transfer Protocol (HTTP).

### Parameters

<i>FrcaHandle</i>	Returns a handle that is required by the other FRCA API subroutines to refer to the newly configured FRCA instance.
<i>Flags</i>	Specifies the behavior of the logging subsystem. The parameter value is constructed by logically ORing single flags. The valid flags are as follows:

**FCTRL\_LOGFORMAT**

**FCTRL\_LOGFORMAT\_ECLF**

**FCTRL\_LOGFORMAT\_VHOST**

**FCTRL\_LOGMODE**

**FCTRL\_LOGMODE\_ON**

## Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

**ENOTREADY** The kernel extension is currently being loaded or unloaded.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlStart** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCtrlStart Subroutine

### Purpose

Starts the interception of TCP data connections for a previously configured Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCtrlStart ( FrcaHandle);
int32_t * FrcaHandle;
```

### Description

The **FrcaCtrlStart** subroutine starts the interception of TCP data connections for an FRCA instance. If the FRCA instance cannot handle the data on that connection, it passes the data to the user-level application that has established the listen socket.

The only FRCA instance type that is currently supported handles static GET requests as part of the Hypertext Transfer Protocol (HTTP).

### Parameters

*FrcaHandle* Identifies the FRCA instance on which this operation is performed.

## Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

ENOENT  
ENOTREADY  
ENOTSOCK

The *FrcaHandle* parameter is invalid.  
The FRCA control instance is in an undefined state.  
A TCP listen socket that was passed in with the *FrcaCtrlCreate* cannot be intercepted because it does not exist.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStop** subroutine.

---

## FrcaCtrlStop Subroutine

### Purpose

Stops the interception of TCP data connections for a Fast Response Cache Accelerator (FRCA) instance.

### Library

FRCA Library (**libfrca.a**)

### Syntax

```
#include <frca.h>
int32_t FrcaCtrlStop ( FrcaHandle);
int32_t * FrcaHandle;
```

### Description

The **FrcaCtrlStop** subroutine stops the interception of newly arriving TCP data connections for a previously configured FRCA instance. Connection requests are passed to the user-level application that has established the listen socket.

The only FRCA instance type that is currently supported handles static GET requests as part of the Hypertext Transfer Protocol (HTTP).

### Parameters

*FrcaHandle*                Identifies the FRCA instance on which this operation is performed.

### Return Values

0	The subroutine completed successfully.
-1	The subroutine failed. The global variable <i>errno</i> is set to indicate the specific type of error.

## Error Codes

ENOENT	The <i>FrcaHandle</i> parameter is invalid.
ENOTREADY	The FRCA control instance has not been started yet.

## Related Information

The **FrcaCacheCreate** subroutine, the **FrcaCacheDelete** subroutine, the **FrcaCacheLoadFile** subroutine, the **FrcaCacheUnloadFile** subroutine, the **FrcaCtrlCreate** subroutine, the **FrcaCtrlDelete** subroutine, the **FrcaCtrlLog** subroutine, the **FrcaCtrlStart** subroutine.

---

## freeaddrinfo Subroutine

### Purpose

Frees memory allocated by the “getaddrinfo Subroutine.”

### Library

The Standard C Library (<**libc.a**>)

### Syntax

```
#include <sys/socket.h>
#include <netdb.h>
void freeaddrinfo (struct addrinfo *ai)
```

### Description

The **freeaddrinfo** subroutine frees one or more **addrinfo** structures returned by the **getaddrinfo** subroutine, along with any additional storage associated with those structures. If the **ai\_next** field of the structure is not NULL, the entire list of structures is freed.

### Parameters

*ai* Points to dynamic storage allocated by the **getaddrinfo** subroutine

## Related Information

“getaddrinfo Subroutine,” and “getnameinfo Subroutine” on page 86.

The **gai\_strerror** Subroutine in *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## getaddrinfo Subroutine

### Purpose

Protocol-independent hostname-to-address translation.

### Library

Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <netdb.h>
int getaddrinfo (hostname, servname, hints, res)
const char *hostname;
const char *servname;
const struct addrinfo *hints;
struct addrinfo **res;
```

## Description

The *hostname* and *servname* parameters describe the hostname and/or service name to be referenced. Zero or one of these arguments may be NULL. A non-NULL hostname may be either a hostname or a numeric host address string (a dotted-decimal for IPv4 or hex for IPv6). A non-NULL servname may be either a service name or a decimal port number.

The *hints* parameter specifies hints concerning the desired return information. The *hostname* and *servname* parameters are pointers to null-terminated strings or NULL. One or both of these arguments must be a non-NULL pointer. In a normal client scenario, both the *hostname* and *servname* parameters are specified. In the normal server scenario, only the *servname* parameter is specified. A non-NULL hostname string can be either a host name or a numeric host address string (for example, a dotted-decimal IPv4 address or an IPv6 hex address). A non-NULL *servname* string can be either a service name or a decimal port number.

The caller can optionally pass an **addrinfo** structure, pointed to by the *hints* parameter, to provide hints concerning the type of socket that the caller supports. In this **hints** structure, all members other than **ai\_flags**, **ai\_family**, **ai\_socktype**, and **ai\_protocol** must be zero or a NULL pointer. A value of PF\_UNSPEC for **ai\_family** means the caller will accept any protocol family. A value of zero for **ai\_socktype** means the caller will accept any socket type. A value of zero for **ai\_protocol** means the caller will accept any protocol. For example, if the caller handles only TCP and not UDP, the **ai\_socktype** member of the **hints** structure should be set to SOCK\_STREAM when the **getaddrinfo** subroutine is called. If the caller handles only IPv4 and not IPv6, the **ai\_family** member of the **hints** structure should be set to PF\_INET when **getaddrinfo** is called. If the *hints* parameter in **getaddrinfo** is a NULL pointer, it is the same as if the caller fills in an **addrinfo** structure initialized to zero with **ai\_family** set to PF\_UNSPEC.

Upon successful return, a pointer to a linked list of one or more **addrinfo** structures is returned through the *res* parameter. The caller can process each **addrinfo** structure in this list by following the **ai\_next** pointer, until a NULL pointer is encountered. In each returned **addrinfo** structure the three members **ai\_family**, **ai\_socktype**, and **ai\_protocol** are the corresponding arguments for a call to the **socket** subroutine. In each **addrinfo** structure, the **ai\_addr** member points to a filled-in socket address structure whose length is specified by the **ai\_addrlen** member.

If the AI\_PASSIVE bit is set in the **ai\_flags** member of the **hints** structure, the caller plans to use the returned socket address structure in a call to the **bind** subroutine. If the *hostname* parameter is a NULL pointer, the IP address portion of the socket address structure will be set to INADDR\_ANY for an IPv4 address or IN6ADDR\_ANY\_INIT for an IPv6 address.

If the AI\_PASSIVE bit is not set in the **ai\_flags** member of the **hints** structure, the returned socket address structure will be ready for a call to the **connect** subroutine (for a connection-oriented protocol) or the **connect**, **sendto**, or **sendmsg** subroutine (for a connectionless protocol). If the *hostname* parameter is a NULL pointer, the IP address portion of the socket address structure will be set to the loopback address.

If the AI\_CANONNAME bit is set in the **ai\_flags** member of the **hints** structure, upon successful return the **ai\_canonname** member of the first **addrinfo** structure in the linked list will point to a NULL-terminated string containing the canonical name of the specified hostname.

If the AI\_NUMERICHOST flag is specified, a non-NULL nodename string supplied is a numeric host address string. Otherwise, an (EAI\_NONAME) error is returned. This flag prevents any type of name resolution service (such as, DNS) from being invoked.

If the AI\_NUMERICSERV flag is specified, a non-NULL servname string supplied is a numeric port string. Otherwise, an (EAI\_NONAME) error is returned. This flag prevents any type of name resolution service (such as, NIS+) from being invoked.

If the `AI_V4MAPPED` flag is specified along with an `ai_family` value of `AF_INET6`, the **getaddrinfo** subroutine returns IPv4-mapped IPv6 addresses when no matching IPv6 addresses (`ai_addrlen` is 16) are found. For example, when using DNS, if no AAAA or A6 records are found, a query is made for A records. Any found are returned as IPv4-mapped IPv6 addresses. The `AI_V4MAPPED` flag is ignored unless `ai_family` equals `AF_INET6`.

If the `AI_ALL` flag is used with the `AI_V4MAPPED` flag, the **getaddrinfo** subroutine returns all matching IPv6 and IPv4 addresses. For example, when using DNS, a query is first made for AAAA/A6 records. If successful, those IPv6 addresses are returned. Another query is made for A records, and any IPv4 addresses found are returned as IPv4-mapped IPv6 addresses. The `AI_ALL` flag without the `AI_V4MAPPED` flag is ignored.

**Note:** When `ai_family` is not specified (`AF_UNSPEC`), `AI_V4MAPPED` and `AI_ALL` flags will only be used if `AF_INET6` is supported.

If the `AI_ADDRCONFIG` flag is specified, a query for AAAA or A6 records should occur only if the node has at least one IPv6 source address configured. A query for A records should occur only if the node has at least one IPv4 source address configured. The loopback address is not considered valid as a configured source address.

All of the information returned by the **getaddrinfo** subroutine is dynamically allocated: the **addrinfo** structures, the socket address structures, and canonical host name strings pointed to by the **addrinfo** structures. To return this information to the system, “freeaddrinfo Subroutine” on page 69 is called.

**The addrinfo structure is defined as:**

```
struct addrinfo {
    int          ai_flags;           /* AI_PASSIVE, AI_CANONNAME */
    int          ai_family;         /* PF_xxx */
    int          ai_socktype;       /* SOCK_xxx */
    int          ai_protocol;       /* 0 or IP=PROTO_xxx for IPv4 and IPv6 */
    size_t       ai_addrlen;        /* length of ai_addr */
    char         *ai_canonname;     /* canonical name for hostname */
    struct sockaddr *ai_addr;       /* binary address */
    struct addrinfo *ai_next;       /* next structure in linked list */
}
```

## Return Values

If the query is successful, a pointer to a linked list of one or more **addrinfo** structures is returned via the *res* parameter. A zero return value indicates success. If the query fails, a non-zero error code will be returned.

## Error Codes

The following names are the non-zero error codes. See *netdb.h* for further definition.

<b>EAI_ADDRFAMILY</b>	Address family for hostname not supported
<b>EAI_AGAIN</b>	Temporary failure in name resolution
<b>EAI_BADFLAGS</b>	Invalid value for <b>ai_flags</b>
<b>EAI_FAIL</b>	Non-recoverable failure in name resolution
<b>EAI_FAMILY</b>	<b>ai_family</b> not supported
<b>EAI_MEMORY</b>	Memory allocation failure
<b>EAI_NODATA</b>	No address associated with <i>hostname</i>
<b>EAI_NONAME</b>	No <i>hostname</i> nor <i>servname</i> provided, or not known
<b>EAI_SERVICE</b>	<i>servname</i> not supported for <b>ai_socktype</b>
<b>EAI_SOCKTYPE</b>	<b>ai_socktype</b> not supported
<b>EAI_SYSTEM</b>	System error returned in <i>errno</i>

## Related Information

“freeaddrinfo Subroutine” on page 69, and “getnameinfo Subroutine” on page 86.

The `gai_strerror` Subroutine in *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 1*.

---

## get\_auth\_method Subroutine

### Purpose

Returns the list of authentication methods for the secure rcmds.

### Library

Authentication Methods Library (**libauthm.a**)

### Syntax

### Description

This method returns the authentication methods currently configured in the order in which they should be attempted in the unsigned integer pointer the user passed in.

The list in the unsigned integer pointer is either NULL (on an error) or is an array of unsigned integers terminated by a zero. Each integer identifies an authentication method. The order that a client should attempt to authenticate is defined by the order of the list.

**Note:** The calling routine is responsible for freeing the memory in which the list is contained.

The flags identifying the authentication methods are defined in the `/usr/include/authm.h` file.

### Parameter

*authm*      Points to an array of unsigned integers. The list of authentication methods is returned in the zero terminated list.

### Return Values

Upon successful completion, the **get\_auth\_method** subroutine returns a zero.

Upon unsuccessful completion, the **get\_auth\_method** subroutine returns an **errno**.

## Related Information

The **chauthent** command, **ftp** command, **lsauthent** command, **rcp** command, **rlogin** command, **rsh** command, **telnet**, **tn**, or **tn3270** command.

The **set\_auth\_method** subroutine.

Communications and networks in *Networks and communication management*.

Authentication and the secure rcmds in *Networks and communication management*.

---

## getdomainname Subroutine

### Purpose

Gets the name of the current domain.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int getdomainname ( Name, NameLen)
char *Name;
int NameLen;
```

### Description

The **getdomainname** subroutine returns the name of the domain for the current processor as previously set by the **setdomainname** subroutine. The returned name is null-terminated unless insufficient space is provided.

The purpose of domains is to enable two distinct networks that may have host names in common to merge. Each network would be distinguished by having a different domain name. Only the Network Information Service (NIS) and the **sendmail** command make use of domains.

All applications containing the **getdomainname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

**Note:** Domain names are restricted to 256 characters.

### Parameters

<i>Name</i>	Specifies the domain name to be returned.
<i>NameLen</i>	Specifies the size of the array pointed to by the <i>Name</i> parameter.

### Return Values

If the call succeeds, a value of 0 is returned. If the call is unsuccessful, a value of -1 is returned and an error code is placed in the **errno** global variable.

### Error Codes

The following error may be returned by this subroutine:

Value	Description
<b>EFAULT</b>	The <i>Name</i> parameter gave an invalid address.

### Related Information

The **gethostname** subroutine, **setdomainname** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## gethostbyaddr Subroutine

### Purpose

Gets network host entry by address.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include <netdb.h>
```

```
struct hostent *gethostbyaddr ( Address, Length, Type)  
const void *Address, size_t Length, int Type;
```

### Description

The **gethostbyaddr** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **gethostbyaddr** subroutine retrieves information about a host using the host address as a search key. Unless specified, the **gethostbyaddr** subroutine uses the default name services ordering, that is, it will query DNS/BIND, NIS, then the local **/etc/hosts** file.

When using DNS/BIND name service resolution, if the file **/etc/resolv.conf** exists, the **gethostbyaddr** subroutine queries the domain name server. The **gethostbyaddr** subroutine recognizes domain name servers as described in RFC 883.

When using NIS for name resolution, if the **getdomainname** subroutine is successful and **yp\_bind** indicates NIS is running, then the **gethostbyaddr** subroutine queries NIS.

The **gethostbyaddr** subroutine also searches the local **/etc/hosts** file when indicated to do so.

The **gethostbyaddr** returns a pointer to a **hostent** structure, which contains information obtained from one of the name resolutions services. The **hostent** structure is defined in the **netdb.h** file.

The environment variable, NSORDER can be set to override the default name services ordering and the order specified in the **/etc/netsvc.conf** file.

### Parameters

<i>Address</i>	Specifies a host address. The host address is passed as a pointer to the binary format address.
<i>Length</i>	Specifies the length of host address.
<i>Type</i>	Specifies the domain type of the host address. It can be either <b>AF_INET</b> or <b>AF_INET6</b> .

### Return Values

The **gethostbyaddr** subroutine returns a pointer to a **hostent** structure upon success.

If an error occurs or if the end of the file is reached, the **gethostbyaddr** subroutine returns a NULL pointer and sets **h\_errno** to indicate the error.

## Error Codes

The **gethostbyaddr** subroutine is unsuccessful if any of the following errors occur:

Error	Description
<b>HOST_NOT_FOUND</b>	The host specified by the <i>Name</i> parameter is not found.
<b>TRY_AGAIN</b>	The local server does not receive a response from an authoritative server. Try again later.
<b>NO_RECOVERY</b>	This error code indicates an unrecoverable error.
<b>NO_ADDRESS</b>	The requested <i>Address</i> parameter is valid but does not have a name at the name server.
<b>SERVICE_UNAVAILABLE</b>	None of the name services specified are running or available.

## Files

<b>/etc/hosts</b>	Contains the host-name database.
<b>/etc/resolv.conf</b>	Contains the name server and domain name information.
<b>/etc/netsvc.conf</b>	Contains the name of the services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

The **endhostent** subroutine, **gethostbyname** subroutine, **sethostent** subroutine, **gethostent** subroutine, **inet\_addr** subroutine.

Sockets Overview, and Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## gethostbyaddr\_r Subroutine

### Purpose

Gets network host entry by address.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include <netdb.h>
int gethostbyaddr_r(Addr, Len, Type, Htent, Ht_data)
const char *Addr, size_t Len, int Type, struct hostent *Htent, struct hostent_data *Ht_data;
```

### Description

This function internally calls the **gethostbyaddr** subroutine and stores the value returned by the **gethostbyaddr** subroutine to the **hostent** structure.

### Parameters

<i>Addr</i>	Points to the host address that is a pointer to the binary format address.
<i>Len</i>	Specifies the length of the address.

<i>Type</i>	Specifies the domain type of the host address. It can be either <b>AF_INET</b> or <b>AF_INET6</b> .
<i>Htent</i>	Points to a <b>hostent</b> structure which is used to store the return value of the <b>gethostaddr</b> subroutine.
<i>Ht_data</i>	Points to a <b>hostent_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** The return value of the **gethostbyaddr** subroutine points to static data that is overwritten by subsequent calls. This data must be copied at every call to be saved for use by subsequent calls. The **gethostbyaddr\_r** subroutine solves this problem.

If the *Name* parameter is a **hostname**, this subroutine searches for a machine with that name as an IP address. Because of this, use the **gethostbyname\_r** subroutine.

## Error Codes

The **gethostbyaddr\_r** subroutine is unsuccessful if any of the following errors occur:

HOST_NOT_FOUND	The host specified by the <i>Name</i> parameter was not found.
TRY_AGAIN	The local server did not receive a response from an authoritative server. Try again later.
NO_RECOVERY	Indicates an unrecoverable error occurred.
NO_ADDRESS	The requested <i>Name</i> parameter is valid but does not have an Internet address at the name server.
SERVICE_UNAVAILABLE	None of the name services specified are running or available.
EINVAL	The <b>hostent</b> pointer is NULL

## Files

<b>/etc/hosts</b>	Contains the host name data base.
<b>/etc/resolv.conf</b>	Contains the name server and domain name.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

“endhostent\_r Subroutine” on page 46, “gethostbyaddr\_r Subroutine” on page 75, “gethostent\_r Subroutine” on page 80, and “sethostent\_r Subroutine” on page 207.

---

## gethostbyname Subroutine

### Purpose

Gets network host entry by name.

### Library

Standard C Library (**libc.a**)  
**(libbind)**  
**(libnis)**  
**(liblocal)**

## Syntax

```
#include <netdb.h>
```

```
struct hostent *gethostbyname ( Name)  
char *Name;
```

## Description

The **gethostbyname** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **gethostbyname** subroutine retrieves host address and name information using a host name as a search key. Unless specified, the **gethostbyname** subroutine uses the default name services ordering, that is, it queries DNS/BIND, NIS or the local **/etc/hosts** file for the name.

When using DNS/BIND name service resolution, if the **/etc/resolv.conf** file exists, the **gethostbyname** subroutine queries the domain name server. The **gethostbyname** subroutine recognizes domain name servers as described in RFC883.

When using NIS for name resolution, if the **getdomaninname** subroutine is successful and **yp\_bind** indicates yellow pages are running, then the **gethostbyname** subroutine queries NIS for the name.

The **gethostbyname** subroutine also searches the local **/etc/hosts** file for the name when indicated to do so.

The **gethostbyname** subroutine returns a pointer to a **hostent** structure, which contains information obtained from a name resolution services. The **hostent** structure is defined in the **netdb.h** header file.

## Parameters

<i>Name</i>	Points to the host name.
-------------	--------------------------

## Return Values

The **gethostbyname** subroutine returns a pointer to a **hostent** structure on success.

If the parameter *Name* passed to **gethostbyname** is actually an IP address, **gethostbyname** will return a non-NULL hostent structure with an IP address as the hostname without actually doing a lookup. Remember to call **inet\_addr** subroutine to make sure *Name* is not an IP address before calling **gethostbyname**. To resolve an IP address call **gethostbyaddr** instead.

If an error occurs or if the end of the file is reached, the **gethostbyname** subroutine returns a null pointer and sets **h\_errno** to indicate the error.

The environment variable, **NSORDER** can be set to override the default name services ordering and the order specified in the **/etc/netsvc.conf** file.

By default, resolver routines first attempt to resolve names through the DNS/BIND, then NIS and the **/etc/hosts** file. The **/etc/netsvc.conf** file may specify a different search order. The environment variable **NSORDER** overrides both the **/etc/netsvc.conf** file and the default ordering. Services are ordered as **hosts = value, value, value** in the **/etc/netsvc.conf** file where at least one value must be specified from the list **bind, nis, local**. **NSORDER** specifies a list of values.

## Error Codes

The **gethostbyname** subroutine is unsuccessful if any of the following errors occur:

Error	Description
<b>HOST_NOT_FOUND</b>	The host specified by the <i>Name</i> parameter was not found.
<b>TRY_AGAIN</b>	The local server did not receive a response from an authoritative server. Try again later.
<b>NO_RECOVERY</b>	This error code indicates an unrecoverable error.
<b>NO_ADDRESS</b>	The requested <i>Name</i> is valid but does not have an Internet address at the name server.
<b>SERVICE_UNAVAILABLE</b>	None of the name services specified are running or available.

## Examples

The following program fragment illustrates the use of the **gethostbyname** subroutine to look up a destination host:

```
hp=gethostbyname(argv[1]);
if(hp == NULL) {
    fprintf(stderr, "rlogin: %s: unknown host\n", argv[1]);
    exit(2);
}
```

## Files

<b>/etc/hosts</b>	Contains the host name data base.
<b>/etc/resolv.conf</b>	Contains the name server and domain name.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

The **endhostent** subroutine, **gethostbyaddr** subroutine, **gethostent** subroutine, **sethostent** subroutine, **inet\_addr** subroutine.

Sockets Overview and Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## gethostbyname\_r Subroutine

### Purpose

Gets network host entry by name.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include netdb.h>
int gethostbyname_r(Name, Htent, Ht_data)

const char *Name, struct hostent *Htent, struct hostent_data *Ht_data;
```

## Description

This function internally calls the **gethostbyname** subroutine and stores the value returned by the **gethostbyname** subroutine to the **hostent** structure.

## Parameters

<i>Name</i>	Points to the host name (which is a constant).
<i>Htent</i>	Points to a <b>hostent</b> structure in which the return value of the <b>gethostbyname</b> subroutine is stored.
<i>Ht_data</i>	Points to a <b>hostent_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

### Note:

The return value of the **gethostbyname** subroutine points to static data that is overwritten by subsequent calls. This data must be copied at every call to be saved for use by subsequent calls. The **gethostbyname\_r** subroutine solves this problem.

If the *Name* parameter is an IP address, this subroutine searches for a machine with that IP address as a name. Because of this, use the **gethostbyaddr\_r** subroutine instead of the **gethostbyname\_r** subroutine if the *Name* parameter is an IP address.

## Error Codes

The **gethostbyname\_r** subroutine is unsuccessful if any of the following errors occurs:

HOST_NOT_FOUND	The host specified by the <i>Name</i> parameter was not found.
TRY_AGAIN	The local server did not receive a response from an authoritative server. Try again later.
NO_RECOVERY	An unrecoverable error occurred.
NO_ADDRESS	The requested <i>Name</i> is valid but does not have an Internet address at the name server.
SERVICE_UNAVAILABLE	None of the name services specified are running or available.
EINVAL	The <b>hostent</b> pointer is NULL.

## Files

<b>/etc/hosts</b>	Contains the host name data base.
<b>/etc/resolv.conf</b>	Contains the name server and domain name.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

“endhostent\_r Subroutine” on page 46, “gethostbyaddr\_r Subroutine” on page 75, “gethostent\_r Subroutine” on page 80, and “sethostent\_r Subroutine” on page 207.

---

## gethostent Subroutine

### Purpose

Retrieves a network host entry.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include <netdb.h>
```

```
struct hostent *gethostent ()
```

### Description

The **gethostent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

When using DNS/BIND name service resolution, the **gethostent** subroutine is not defined.

When using NIS name service resolution or searching the local **/etc/hosts** file, the **gethostent** subroutine reads the next line of the **/etc/hosts** file, opening the file if necessary.

The **gethostent** subroutine returns a pointer to a **hostent** structure, which contains the equivalent fields for a host description line in the **/etc/hosts** file. The **hostent** structure is defined in the **netdb.h** file.

### Return Values

Upon successful completion, the **gethostent** subroutine returns a pointer to a **hostent** structure.

If an error occurs or the end of the file is reached, the **gethostent** subroutine returns a null pointer.

### Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

### Related Information

The **gethostbyaddr** subroutine, **gethostbyname** subroutine, **sethostent** subroutine **endhostent** subroutine.

Sockets Overview and Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## gethostent\_r Subroutine

### Purpose

Retrieves a network host entry.

## Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

## Syntax

```
#include <netdb.h>
```

```
int gethostent_r (htent, ht_data)
struct hostent *htent;
struct hostent_data *ht_data;
```

## Description

When using DNS/BIND name service resolution, the **gethostent\_r** subroutine is not defined.

When using NIS name service resolution or searching the local **/etc/hosts** file, the **gethostent\_r** subroutine reads the next line of the **/etc/hosts** file, and opens the file if necessary.

The **gethostent\_r** subroutine internally calls the **gethostent** subroutine, and stores the values in the **htent** and **ht\_data** structures.

The **gethostent** subroutine overwrites the static data returned in subsequent calls. The **gethostent\_r** subroutine does not.

## Parameters

<i>htent</i>	Points to the <b>hostent</b> structure
<i>ht_data</i>	Points to the <b>hostent_data</b> structure

## Return Values

This subroutine returns a 0 if successful, and a -1 if unsuccessful.

## Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/usr/include/netdb.h</b>	Contains the network database structure.

## Related Information

“gethostbyaddr\_r Subroutine” on page 75, “gethostbyname\_r Subroutine” on page 78, “sethostent\_r Subroutine” on page 207, and “endhostent\_r Subroutine” on page 46.

---

## gethostid Subroutine

### Purpose

Gets the unique identifier of the current host.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int gethostid ( )
```

## Description

The **gethostid** subroutine allows a process to retrieve the 32-bit identifier for the current host. In most cases, the host ID is stored in network standard byte order and is a DARPA Internet address for a local machine.

All applications containing the **gethostid** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Return Values

Upon successful completion, the **gethostid** subroutine returns the identifier for the current host.

## Related Information

The **gethostname** subroutine, **sethostid** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## gethostname Subroutine

### Purpose

Gets the name of the local host.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <unistd.h>
int gethostname ( Name,  NameLength)
char *Name;
size_t NameLength;
```

### Description

The **gethostname** subroutine retrieves the standard host name of the local host. If excess space is provided, the returned *Name* parameter is null-terminated. If insufficient space is provided, the returned name is truncated to fit in the given space. System host names are limited to 256 characters.

The **gethostname** subroutine allows a calling process to determine the internal host name for a machine on a network.

All applications containing the **gethostname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

*Name*                      Specifies the address of an array of bytes where the host name is to be stored.

*NameLength* Specifies the length of the *Name* array.

## Return Values

Upon successful completion, the system returns a value of 0.

If the **gethostname** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **gethostname** subroutine is unsuccessful if the following is true:

Error	Description
<b>EFAULT</b>	The <i>Name</i> parameter or <i>NameLength</i> parameter gives an invalid address.

## Related Information

The **gethostid** subroutine, **sethostid** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## GetMultipleCompletionStatus Subroutine

### Purpose

Dequeues multiple completion packets from a specified I/O completion port.

### Syntax

```
#include <sys/iocp.h>

int GetMultipleCompletionStatus (CompletionPort, Nmin, Nmax, Timeout, Results[])
HANDLE CompletionPort;
DWORD Nmin, Nmax, Timeout;
struct gmcs {
    DWORD transfer_count, completion_key, errno;
    LPOVERLAPPED overlapped;
} Results[];
```

### Description

The **GetMultipleCompletionStatus** subroutine attempts to dequeue a number of completion packets from the completion port that is specified by the *CompletionPort* parameter. The number of dequeued completion packets that are wanted ranges from the value of the *Nmin* parameter through the value of the *Nmax* parameter. As it collects the packets, this subroutine might wait a predetermined maximum amount of time that is specified by the *Timeout* parameter for the minimum number of completion packets to arrive. If, for example, the X<sup>th</sup> completion packet does not arrive in time, the subroutine returns with only X-1 packets completed.

Either the *Timeout* parameter or a signal might cause a return with completions fewer than the value of the *Nmin* parameter. In other words, *Nmin* completions are not guaranteed to be returned unless the *Timeout* parameter value is set to INFINITE, and a signal does not interrupt the wait. The return of zero completions is not considered an error. The **errno** value will, however, indicate the condition with either the **ETIMEDOUT** or **EINTR** error code. In extreme low-memory situations, the kernel might not be able to provide a timeout. In this case, the system call returns immediately with any available completions, up to the value of the *Nmax* parameter, and the **errno** value is set to **ENOMEM**. Be sure to set the **errno** value

to zero before calling the **GetMultipleCompletionStatus** subroutine so that the change of the **errno** value that the subroutine makes can be distinguished from the existing value.

The **GetMultipleCompletionStatus** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

#### Notes:

1. This subroutine only works with file descriptors of sockets, or regular files for use with the asynchronous I/O subsystem. It does not work with file descriptors of other types.
2. This function must be the exclusive wait mechanism on a completion port. Multiple simultaneous waits through the **GetMultipleCompletionStatus** subroutine, the **GetQueuedCompletionStatus** subroutine, or both, are not supported.
3. When the **GetMultipleCompletionStatus** subroutine is used with the **lio\_listio** subroutine, you can set the value of the *cmd* parameter of the **lio\_listio** subroutine to **LIO\_NOWAIT\_GMCS** to avoid asynchronous updating of the **aiocb** structures, thereby reducing overhead. In this case, you must use the **GetMultipleCompletionStatus** subroutine to wait for I/O completions, and retrieve completion status only from the **struct gmcs** members, not from the **aiocb** structure. When using the **LIO\_NOWAIT\_GMCS** value, do not use the *completion\_key* value in the **gmcs** structure. Do not use the **LIO\_NOWAIT\_AIOWAIT** value with the **lio\_listio** subroutine when using the **GetMultipleCompletionStatus** subroutine. The **LIO\_NOWAIT\_GMCS** value is available for that purpose.
4. Cancelling an asynchronous I/O operation will not affect the **GetMultipleCompletionStatus** subroutine. Even if the cancelling reduces the number of active asynchronous I/O operations to zero, the subroutine will continue to wait.
5. When using the **GetMultipleCompletionStatus** subroutine with sockets, do not wait for multiple completions (*Nmin* > 1) with an **INFINITE** timeout. Use a finite timeout value, and to be prepared to repeat the call if additional completions are still expected.

## Parameters

<i>CompletionPort</i>	Specifies the file descriptor for the completion port that this subroutine will access.
<i>Nmin</i>	Specifies the minimum number of completions. Fewer might be returned if the value of the <i>timeout</i> parameter is exceeded, or a signal accepted. More might be returned, up to the number that is specified by the <i>Nmax</i> parameter, if additional completions have occurred. Setting the value of the <i>Nmin</i> parameter to zero will poll for completions and return immediately, ignoring the value of the <i>timeout</i> parameter.
<i>Nmax</i>	Specifies the maximum number of completions to wait for, up to the value of the <b>GMCS_NMAX</b> macro.
<i>Results</i>	This is the address of an array of the <b>gmcs</b> structure to receive the completion data. The array must contain space for the number of entries specified by the <i>Nmax</i> parameter.
<i>Results[i].transfer_count</i>	Specifies the number of bytes transferred. This parameter is set by the subroutine from the value received in the <i>i</i> <sup>th</sup> completion packet. This value is limited to 2 G.
<i>Results[i].completion_key</i>	Specifies the completion key associated with the file descriptor that is used in the transfer request. This parameter is set by the subroutine from the value received in the <i>i</i> <sup>th</sup> completion packet. Do not use this value with the <b>LIO_NOWAIT_GMCS</b> command parameter of the <b>lio_listio</b> subroutine.
<i>Results[i].errno</i>	Specifies the <b>errno</b> value that is associated with the <i>i</i> <sup>th</sup> completion packet. When asynchronous I/O requests are started using the <b>lio_listio</b> subroutine with the <b>LIO_NOWAIT_GMCS</b> command parameter, you must use this error value, not the <b>aio_errno</b> member in the <b>aiocb</b> structure, to retrieve the error value that is associated with an I/O request.

<i>Results[i].overlapped</i>	Specifies the overlapped structure that is used in the transfer request. This parameter is set by the subroutine from the value received in the <i>i</i> <sup>th</sup> completion packet. For regular files, this parameter contains a pointer to the asynchronous I/O control block (AIOCB) for a completed AIO request. If an application uses the same completion port for both socket and AIO to regular files, it must use unique <i>completion_key</i> values to differentiate between sockets and regular files to properly interpret the <i>overlapped</i> parameter.
<i>Timeout</i>	Specifies the amount of time in milliseconds that the subroutine is to wait for completion packets. This value can be set to zero. If this parameter is set to INFINITE, the subroutine will never time out.

## Return Values

Success	The subroutine returns an integer ranging from zero through the value of the <i>Nmax</i> parameter, indicating how many completion packets are dequeued.
Failure	The subroutine returns a value of -1.

## Error codes

The subroutine is unsuccessful if any of the following errors occur:

<b>EINVAL</b>	The value of the <i>CompletionPort</i> or other parameter is not valid.
<b>EBUSY</b>	Another thread is already waiting on the I/O completion port.
<b>EBADF</b>	This error code might also be returned when the value of the <i>CompletionPort</i> parameter is not valid.

If an error occurs after some completions have been handled, the error notifications will be lost. An **EFAULT** error when copying out results can cause the situation.

## Examples

1. The following program fragment illustrates the use of the **GetMultipleCompletionStatus** subroutine to dequeue up to 10 completion packets within a 100-millisecond window.

```
struct gmcs results[10];
int n_results;
HANDLE iocpfd;
errno = 0;
n_results = GetMultipleCompletionStatus(iocpfd, 10, 10, 100, results);
```

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “ReadFile Subroutine” on page 165, “WriteFile Subroutine” on page 258, “CreateloCompletionPort Subroutine” on page 37, “GetQueuedCompletionStatus Subroutine” on page 101, and “PostQueuedCompletionStatus Subroutine” on page 161.

The **lio\_listio** subroutine in *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 1*.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## getnameinfo Subroutine

### Purpose

Address-to-host name translation [given the binary address and port].

**Note:** This is the reverse functionality of the “getaddrinfo Subroutine” on page 69 host-to-address translation.

**Attention:** This is not a POSIX (1003.1g) specified function.

### Library

Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <netdb.h>
int
getnameinfo (sa, salen, host, hostlen, serv, servlen, flags)
const struct sockaddr *sa;
char *host;
size_t hostlen;
char *serv;
size_t servlen;
int flags;
```

### Description

The *sa* parameter points to either a **sockaddr\_in** structure (for IPv4) or a **sockaddr\_in6** structure (for IPv6) that holds the IP address and port number. The *salen* parameter gives the length of the **sockaddr\_in** or **sockaddr\_in6** structure.

**Note:** A reverse lookup is performed on the IP address and port number provided in *sa*.

The *host* parameter is a buffer where the hostname associated with the IP address is copied. The *hostlen* parameter provides the length of this buffer. The service name associated with the port number is copied into the buffer pointed to by the *serv* parameter. The *servlen* parameter provides the length of this buffer.

The *flags* parameter defines flags that may be used to modify the default actions of this function. By default, the fully-qualified domain name (FQDN) for the host is looked up in DNS and returned.

<b>NI_NOFQDN</b>	If set, return only the hostname portion of the FQDN. If cleared, return the FQDN.
<b>NI_NUMERICHOST</b>	If set, return the numeric form of the host address. If cleared, return the name.
<b>NI_NAMEREQD</b>	If set, return an error if the host's name cannot be determined. If cleared, return the numeric form of the host's address (as if <b>NI_NUMERICHOST</b> had been set).
<b>NI_NUMERICSERV</b>	If set, return the numeric form of the desired service. If cleared, return the service name.
<b>NI_DGRAM</b>	If set, consider the desired service to be a datagram service, (for example, call <b>getservbyport</b> with an argument of <b>udp</b> ). If clear, consider the desired service to be a stream service (for example, call <b>getservbyport</b> with an argument of <b>tcp</b> ).

## Return Values

A zero return value indicates successful completion; a non-zero value indicates failure. If successful, the strings for hostname and service name are copied into the *host* and *serv* buffers, respectively. If either the host or service name cannot be located, the numeric form is copied into the *host* and *serv* buffers, respectively.

## Related Information

“getaddrinfo Subroutine” on page 69, and “freeaddrinfo Subroutine” on page 69.

The `gai_strerror` Subroutine in *AIX Version 6.1 Technical Reference: Base Operating System and Extensions Volume 1*.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## getnetbyaddr Subroutine

### Purpose

Gets network entry by address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct netent *getnetbyaddr (Network, Type)
long Network;
int Type;
```

### Description

The **getnetbyaddr** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getnetbyaddr** subroutine retrieves information from the **/etc/networks** file using the network address as a search key. The **getnetbyaddr** subroutine searches the file sequentially from the start of the file until it encounters a matching net number and type or until it reaches the end of the file.

The **getnetbyaddr** subroutine returns a pointer to a **netent** structure, which contains the equivalent fields for a network description line in the **/etc/networks** file. The **netent** structure is defined in the **netdb.h** file.

Use the **endnetent** subroutine to close the **/etc/networks** file.

All applications containing the **getnetbyaddr** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Network</i>	Specifies the number of the network to be located.
<i>Type</i>	Specifies the address family for the network. The only supported value is <b>AF_INET</b> .

## Return Values

Upon successful completion, the **getnetbyaddr** subroutine returns a pointer to a **netent** structure.

If an error occurs or the end of the file is reached, the **getnetbyaddr** subroutine returns a null pointer.

## Files

**/etc/networks**

Contains official network names.

## Related Information

The **endnetent** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **setnetent** subroutine.

Sockets Overview in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## getnetbyaddr\_r Subroutine

### Purpose

Gets network entry by address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include<netdb.h>
int getnetbyaddr_r(net, type, netent, net_data)

register in_addr_t net;
register int type;
struct netent *netent;
struct netent_data *net_data;
```

### Description

The **getnetbyaddr\_r** subroutine retrieves information from the **/etc/networks** file using the *Name* parameter as a search key.

The **getnetbyaddr\_r** subroutine internally calls the **getnetbyaddr** subroutine and stores the information in the structure data.

The **getnetbyaddr** subroutine overwrites the static data returned in subsequent calls. The **getnetbyaddr\_r** subroutine does not.

Use the **endnetent\_r** subroutine to close the **/etc/networks** file.

### Parameters

*Net*

Specifies the number of the network to be located.

*Type*

Specifies the address family for the network. The only supported values are AF\_INET, and AF\_INET6.

*netent*

Points to the **netent** structure.

*net\_data*

Points to the **net\_data** structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

## Files

**/etc/networks**

Contains official network names.

## Related Information

“endnetent\_r Subroutine” on page 47, “getnetbyname\_r Subroutine” on page 90, “getnetent\_r Subroutine” on page 92, and “setnetent\_r Subroutine” on page 211.

---

## getnetbyname Subroutine

### Purpose

Gets network entry by name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct netent *getnetbyname (Name)  
char *Name;
```

### Description

The **getnetbyname** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getnetbyname** subroutine retrieves information from the **/etc/networks** file using the *Name* parameter as a search key. The **getnetbyname** subroutine searches the **/etc/networks** file sequentially from the start of the file until it encounters a matching net name or until it reaches the end of the file.

The **getnetbyname** subroutine returns a pointer to a **netent** structure, which contains the equivalent fields for a network description line in the **/etc/networks** file. The **netent** structure is defined in the **netdb.h** file.

Use the **endnetent** subroutine to close the **/etc/networks** file.

All applications containing the **getnetbyname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

*Name*                      Points to a string containing the name of the network.

## Return Values

Upon successful completion, the **getnetbyname** subroutine returns a pointer to a **netent** structure.

If an error occurs or the end of the file is reached, the **getnetbyname** subroutine returns a null pointer.

## Files

**/etc/networks**

Contains official network names.

## Related Information

The **endnetent** subroutine, **getnetbyaddr** subroutine, **getnetent** subroutine, **setnetent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getnetbyname\_r Subroutine

### Purpose

Gets network entry by name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int getnetbyname_r(Name, netent, net_data)
register const char *Name;
struct netent *netent;
struct netent_data *net_data;
```

### Description

The **getnetbyname\_r** subroutine retrieves information from the **/etc/networks** file using the *Name* parameter as a search key.

The **getnetbyname\_r** subroutine internally calls the **getnetbyname** subroutine and stores the information in the structure data.

The **getnetbyname** subroutine overwrites the static data returned in subsequent calls. The **getnetbyname\_r** subroutine does not.

Use the **endnetent\_r** subroutine to close the **/etc/networks** file.

### Parameters

<i>Name</i>	Points to a string containing the name of the network.
<i>netent</i>	Points to the <b>netent</b> structure.
<i>net_data</i>	Points to the <b>net_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getnetbyname\_r** subroutine returns a -1 to indicate error.

## Files

**/etc/networks**

Contains official network names.

## Related Information

“endnetent\_r Subroutine” on page 47, “getnetbyaddr\_r Subroutine” on page 88, “getnetent\_r Subroutine” on page 92, and “setnetent\_r Subroutine” on page 211.

---

## getnetent Subroutine

### Purpose

Gets network entry.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
struct netent *getnetent ( )
```

### Description

The **getnetent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getnetent** subroutine retrieves network information by opening and sequentially reading the **/etc/networks** file.

The **getnetent** subroutine returns a pointer to a **netent** structure, which contains the equivalent fields for a network description line in the **/etc/networks** file. The **netent** structure is defined in the **netdb.h** file.

Use the **endnetent** subroutine to close the **/etc/networks** file.

All applications containing the **getnetent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Return Values

Upon successful completion, the **getnetent** subroutine returns a pointer to a **netent** structure.

If an error occurs or the end of the file is reached, the **getnetent** subroutine returns a null pointer.

## Files

**/etc/networks**

Contains official network names.

## Related Information

The **endnetent** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **setnetent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getnetent\_r Subroutine

### Purpose

Gets network entry.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int getnetent_r(netent, net_data)
```

```
struct netent *netent;
struct netent_data *net_data;
```

### Description

The **getnetent\_r** subroutine retrieves network information by opening and sequentially reading the **/etc/networks** file. This subroutine internally calls the **getnetent** subroutine and stores the values in the hostent structure.

The **getnetent** subroutine overwrites the static data returned in subsequent calls. The **getnetent\_r** subroutine does not. Use the **endnetent\_r** subroutine to close the **/etc/networks** file.

### Parameters

<i>netent</i>	Points to the <b>netent</b> structure.
<i>net_data</i>	Points to the <b>net_data</b> structure.

### Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getnetent\_r** subroutine returns a -1 to indicate error.

### Files

<b>/etc/networks</b>	Contains official network names.
----------------------	----------------------------------

## Related Information

“endnetent\_r Subroutine” on page 47, “getnetbyaddr\_r Subroutine” on page 88, “getnetbyname\_r Subroutine” on page 90, and “setnetent\_r Subroutine” on page 211.

---

## getnetgrent\_r Subroutine

### Purpose

Handles the group network entries.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include<netdb.h>
int getnetgrent_r(machinep, namep, domainp, ptr)
    char **machinep, **namep, **domainp;
void **ptr;
```

### Description

The **getnetgrent\_r** subroutine internally calls the **getnetgrent** subroutine and stores the information in the structure data. This subroutine returns 1 or 0, depending if netgroup contains the machine, user, and domain triple as a member. Any of these three strings can be NULL, in which case it signifies a wild card.

The **getnetgrent\_r** subroutine returns the next member of a network group. After the call, the *machinep* parameter contains a pointer to a string containing the name of the machine part of the network group member. The *namep* and *domainp* parameters contain similar pointers. If *machinep*, *namep*, or *domainp* is returned as a NULL pointer, it signifies a wild card.

The **getnetgrent** subroutine overwrites the static data returned in subsequent calls. The **getnetgrent\_r** subroutine does not.

### Parameters

<i>machinep</i>	Points to the string containing the machine part of the network group.
<i>namep</i>	Points to the string containing the user part of the network group.
<i>domainp</i>	Points to the string containing the domain name.
<i>ptr</i>	Keeps the function threadsafe.

### Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

### Files

<b>/etc/netgroup</b>	Contains network groups recognized by the system.
<b>/usr/include/netdb.h</b>	Contains the network database structures.

### Related Information

“endnetgrent\_r Subroutine” on page 48, and “setnetgrent\_r Subroutine” on page 211.

---

## getpeername Subroutine

### Purpose

Gets the name of the peer socket.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
int getpeername ( Socket, Name, NameLength)
int Socket;
struct sockaddr *Name;
socklen_t *NameLength;
```

### Description

The **getpeername** subroutine retrieves the *Name* parameter from the peer socket connected to the specified socket. The *Name* parameter contains the address of the peer socket upon successful completion.

A process created by another process can inherit open sockets. The created process may need to identify the addresses of the sockets it has inherited. The **getpeername** subroutine allows a process to retrieve the address of the peer socket at the remote end of the socket connection.

**Note:** The **getpeername** subroutine operates only on connected sockets.

A process can use the **getsockname** subroutine to retrieve the local address of a socket.

All applications containing the **getpeername** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the descriptor number of a connected socket.
<i>Name</i>	Points to a <b>sockaddr</b> structure that contains the address of the destination socket upon successful completion. The <b>/usr/include/sys/socket.h</b> file defines the <b>sockaddr</b> structure.
<i>NameLength</i>	Points to the size of the address structure. Initializes the <i>NameLength</i> parameter to indicate the amount of space pointed to by the <i>Name</i> parameter. Upon successful completion, it returns the actual size of the <i>Name</i> parameter returned.

### Return Values

Upon successful completion, a value of 0 is returned and the *Name* parameter holds the address of the peer socket.

If the **getpeername** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **getpeername** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EINVAL</b>	The socket has been shut down.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ENOTCONN</b>	The socket is not connected.
<b>ENOBUFS</b>	Insufficient resources were available in the system to complete the call.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.

## Examples

The following program fragment illustrates the use of the **getpeername** subroutine to return the address of the peer connected on the other end of the socket:

```
struct sockaddr_in name;
int namelen = sizeof(name);
.
.
.
if(getpeername(0,(struct sockaddr*)&name, &namelen)<0){
    syslog(LOG_ERR,"getpeername: %m");
    exit(1);
} else
    syslog(LOG_INFO,"Connection from %s",inet_ntoa(name.sin_addr));
.
.
.
```

## Related Information

The **accept** subroutine, **bind** subroutine, **getsockname** subroutine, **socket** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getprotobyname Subroutine

### Purpose

Gets protocol entry from the */etc/protocols* file by protocol name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct protoent *getprotobyname (Name)
char *Name;
```

### Description

The **getprotobyname** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getprotobyname** subroutine retrieves protocol information from the **/etc/protocols** file by protocol name. An application program can use the **getprotobyname** subroutine to access a protocol name, its aliases, and protocol number.

The **getprotobyname** subroutine searches the **protocols** file sequentially from the start of the file until it finds a matching protocol name or until it reaches the end of the file. The subroutine returns a pointer to a **protoent** structure, which contains fields for a line of information in the **/etc/protocols** file. The **netdb.h** file defines the **protoent** structure.

Use the **endprotoent** subroutine to close the **/etc/protocols** file.

All applications containing the **getprotobyname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*Name*                      Specifies the protocol name.

## Return Values

Upon successful completion, the **getprotobyname** subroutine returns a pointer to a **protoent** structure.

If an error occurs or the end of the file is reached, the **getprotobyname** subroutine returns a null pointer.

## Related Information

The **endprotoent** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **setprotoent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getprotobyname\_r Subroutine

### Purpose

Gets protocol entry from the **/etc/protocols** file by protocol name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>

int getprotobyname_r(Name, protoent, proto_data)
register const char *Name;
struct protoent *protoent;
struct protoent_data *proto_data;
```

### Description

The **getprotobyname\_r** subroutine retrieves protocol information from the **/etc/protocols** file by protocol name.

An application program can use the **getprotobyname\_r** subroutine to access a protocol name, aliases, and protocol number.

The **getprotobyname\_r** subroutine searches the protocols file sequentially from the start of the file until it finds a matching protocol name or until it reaches the end of the file. The subroutine writes the protoent structure, which contains fields for a line of information in the **/etc/protocols** file.

The **netdb.h** file defines the protoent structure.

The **getprotobyname** subroutine overwrites any static data returned in subsequent calls. The **getprotobyname\_r** subroutine does not.

Use the **endprotoent\_r** subroutine to close the **/etc/protocols** file.

## Parameters

<i>Name</i>	Specifies the protocol name.
<i>protoent</i>	Points to the <b>protoent</b> structure.
<i>proto_data</i>	Points to the <b>proto_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getprotobyname\_r** subroutine returns a -1 to indicate error.

## Related Information

“endprotoent\_r Subroutine” on page 50, “getprotobynumber\_r Subroutine” on page 98, “getprotoent\_r Subroutine” on page 100, and “setprotoent\_r Subroutine” on page 213.

---

## getprotobynumber Subroutine

### Purpose

Gets a protocol entry from the **/etc/protocols** file by number.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct protoent *getprotobynumber ( Protocol )
int Protocol;
```

### Description

The **getprotobynumber** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getprotobynumber** subroutine retrieves protocol information from the **/etc/protocols** file using a specified protocol number as a search key. An application program can use the **getprotobynumber** subroutine to access a protocol name, its aliases, and protocol number.

The **getprotobynumber** subroutine searches the **/etc/protocols** file sequentially from the start of the file until it finds a matching protocol name or protocol number, or until it reaches the end of the file. The subroutine returns a pointer to a **protoent** structure, which contains fields for a line of information in the **/etc/protocols** file. The **netdb.h** file defines the **protoent** structure.

Use the **endprotoent** subroutine to close the **/etc/protocols** file.

All applications containing the **getprotobynumber** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Protocol</i>	Specifies the protocol number.
-----------------	--------------------------------

## Return Values

Upon successful completion, the **getprotobynumber** subroutine, returns a pointer to a **protoent** structure.

If an error occurs or the end of the file is reached, the **getprotobynumber** subroutine returns a null pointer.

## Files

<b>/etc/protocols</b>	Contains protocol information.
-----------------------	--------------------------------

## Related Information

The **endprotoent** subroutine, **getprotobyname** subroutine, **getprotoent** subroutine, **setprotoent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getprotobynumber\_r Subroutine

### Purpose

Gets a protocol entry from the **/etc/protocols** file by number.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int getprotobynumber_r(proto, protoent, proto_data)
register int proto;
struct protoent *protoent;
struct protoent_data *proto_data;
```

### Description

The **getprotobynumber\_r** subroutine retrieves protocol information from the **/etc/protocols** file using a specified protocol number as a search key.

An application program can use the **getprotobynumber\_r** subroutine to access a protocol name, aliases, and number.

The **getprotobynumber\_r** subroutine searches the **/etc/protocols** file sequentially from the start of the file until it finds a matching protocol name, protocol number, or until it reaches the end of the file.

The subroutine writes the **protoent** structure, which contains fields for a line of information in the **/etc/protocols** file.

The **netdb.h** file defines the **protoent** structure.

The **getprotobynumber** subroutine overwrites static data returned in subsequent calls. The **getprotobynumber\_r** subroutine does not.

Use the **endprotoent\_r** subroutine to close the **/etc/protocols** file.

## Parameters

<i>proto</i>	Specifies the protocol number.
<i>protoent</i>	Points to the <b>protoent</b> structure.
<i>proto_data</i>	Points to the <b>proto_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getprotobynumber\_r** subroutine sets the *protoent* parameter to NULL and returns a -1 to indicate error.

## Files

<b>/etc/protocols</b>	Contains protocol information.
-----------------------	--------------------------------

## Related Information

“endprotoent\_r Subroutine” on page 50, “getprotobyname\_r Subroutine” on page 96, “getprotoent\_r Subroutine” on page 100, and “setprotoent\_r Subroutine” on page 213.

---

## getprotoent Subroutine

### Purpose

Gets protocol entry from the **/etc/protocols** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
struct protoent *getprotoent ( )
```

## Description

The **getprotoent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getprotoent** subroutine retrieves protocol information from the **/etc/protocols** file. An application program can use the **getprotoent** subroutine to access a protocol name, its aliases, and protocol number.

The **getprotoent** subroutine opens and performs a sequential read of the **/etc/protocols** file. The **getprotoent** subroutine returns a pointer to a **protoent** structure, which contains the fields for a line of information in the **/etc/protocols** file. The **netdb.h** file defines the **protoent** structure.

Use the **endprotoent** subroutine to close the **/etc/protocols** file.

All applications containing the **getprotoent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Return Values

Upon successful completion, the **getprotoent** subroutine returns a pointer to a **protoent** structure.

If an error occurs or the end of the file is reached, the **getprotoent** subroutine returns a null pointer.

## Files

<b>/etc/protocols</b>	Contains protocol information.
-----------------------	--------------------------------

## Related Information

The **endprotoent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **setprotoent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getprotoent\_r Subroutine

### Purpose

Gets protocol entry from the **/etc/protocols** file.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>

int getprotoent_r(protoent, proto_data)
struct protoent *protoent;
struct protoent_data *proto_data;
```

### Description

The **getprotoent\_r** subroutine retrieves protocol information from the **/etc/protocols** file. An application program can use the **getprotoent\_r** subroutine to access a protocol name, its aliases, and protocol

number. The **getprotoent\_r** subroutine opens and performs a sequential read of the **/etc/protocols** file. This subroutine writes to the **protoent** structure, which contains the fields for a line of information in the **/etc/protocols** file.

The **netdb.h** file defines the **protoent** structure.

Use the **endprotoent\_r** subroutine to close the **/etc/protocols** file. Static data is overwritten in subsequent calls when using the **getprotoent** subroutine. The **getprotoent\_r** subroutine does not overwrite.

## Parameters

<i>protoent</i>	Points to the <b>protoent</b> structure.
<i>proto_data</i>	Points to the <b>proto_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getprotoent\_r** subroutine sets the *protoent* parameter to NULL.

## Files

<b>/etc/protocols</b>	Contains protocol information.
-----------------------	--------------------------------

## Related Information

“endprotoent\_r Subroutine” on page 50, “setprotoent\_r Subroutine” on page 213, “getprotobyname\_r Subroutine” on page 96, “getprotobyname\_r Subroutine” on page 98, and “setprotoent\_r Subroutine” on page 213.

---

## GetQueuedCompletionStatus Subroutine

### Purpose

Dequeues a completion packet from a specified I/O completion port.

### Syntax

```
#include <iocp.h>
boolean_t GetQueuedCompletionStatus (CompletionPort, TransferCount, CompletionKey, Overlapped, Timeout)
HANDLE CompletionPort;
LPDWORD TransferCount, CompletionKey;
LPOVERLAPPED Overlapped; DWORD Timeout;
```

### Description

The **GetQueuedCompletionStatus** subroutine attempts to dequeue a completion packet from the *CompletionPort* parameter. If there is no completion packet to be dequeued, this subroutine waits a predetermined amount of time as indicated by the *Timeout* parameter for a completion packet to arrive.

The **GetQueuedCompletionStatus** subroutine returns a boolean indicating whether or not a completion packet has been dequeued.

The **GetQueuedCompletionStatus** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

**Note:** This subroutine only works with file descriptors of sockets, or regular files for use with the Asynchronous I/O (AIO) subsystem. It does not work with file descriptors of other types.

## Parameters

<i>CompletionPort</i>	Specifies the completion port that this subroutine will attempt to access.
<i>TransferCount</i>	Specifies the number of bytes transferred. This parameter is set by the subroutine from the value received in the completion packet.
<i>CompletionKey</i>	Specifies the completion key associated with the file descriptor used in the transfer request. This parameter is set by the subroutine from the value received in the completion packet.
<i>Overlapped</i>	Specifies the overlapped structure used in the transfer request. This parameter is set by the subroutine from the value received in the completion packet. For regular files, this parameter contains a pointer to the AIOCB for a completed AIO request. If an application uses the same completion port for both socket and AIO to regular files, it must use unique <i>CompletionKey</i> values to differentiate between sockets and regular files in order to properly interpret the <i>Overlapped</i> parameter.
<i>Timeout</i>	Specifies the amount of time in milliseconds the subroutine is to wait for a completion packet. If this parameter is set to INFINITE, the subroutine will never timeout.

## Return Values

Upon successful completion, the **GetQueuedCompletionStatus** subroutine returns a boolean indicating its success.

If the **GetQueuedCompletionStatus** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of 0 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable, see the link in the Related Information section of this document.

## Error Codes

The subroutine is unsuccessful if any of the following errors occur:

<b>ETIMEDOUT</b>	No completion packet arrived to be dequeued and the <i>Timeout</i> parameter has elapsed.
<b>EINVAL</b>	The value of the <i>CompletionPort</i> or other parameter is not valid.
<b>EAGAIN</b>	Resource temporarily unavailable. If a sleep is interrupted by a signal, <b>EAGAIN</b> may be returned.
<b>ENOTCONN</b>	Socket is not connected. The <b>ENOTCONN</b> return can happen for two reasons. One is if a request is made, the <b>fd</b> is then closed, then the request is returned back to the process. The error will be <b>ENOTCONN</b> . The other is if the socket drops while the <b>fd</b> is still open, the requests after the socket drops (disconnects) will return <b>ENOTCONN</b> .
<b>EBADF</b>	This error code might also be returned when the value of the <i>CompletionPort</i> parameter is not valid.

## Examples

The following program fragment illustrates the use of the **GetQueuedCompletionStatus** subroutine to dequeue a completion packet.

```
int transfer_count, completion_key
LPOVERLAPPED overlapped;
c = GetQueuedCompletionStatus (34, &transfer_count, &completion_key, &overlapped, 1000);
```

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “ReadFile Subroutine” on page 165, “WriteFile Subroutine” on page 258, “CreateloCompletionPort Subroutine” on page 37, “GetMultipleCompletionStatus Subroutine” on page 83, and “PostQueuedCompletionStatus Subroutine” on page 161.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## getservbyname Subroutine

### Purpose

Gets service entry by name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct servent *getservbyname ( Name, Protocol)
char *Name, *Protocol;
```

### Description

The **getservbyname** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getservbyname** subroutine retrieves an entry from the **/etc/services** file using the service name as a search key.

An application program can use the **getservbyname** subroutine to access a service, service aliases, the protocol for the service, and a protocol port number for the service.

The **getservbyname** subroutine searches the **/etc/services** file sequentially from the start of the file until it finds one of the following:

- Matching name and protocol number
- Matching name when the *Protocol* parameter is set to 0
- End of the file

Upon locating a matching name and protocol, the **getservbyname** subroutine returns a pointer to the **servent** structure, which contains fields for a line of information from the **/etc/services** file. The **netdb.h** file defines the **servent** structure and structure fields.

Use the **endservent** subroutine to close the **/etc/services** file.

All applications containing the **getservbyname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Name</i>	Specifies the name of a service.
<i>Protocol</i>	Specifies a protocol for use with the specified service.

## Return Values

The **getservbyname** subroutine returns a pointer to a **servent** structure when a successful match occurs. Entries in this structure are in network byte order.

If an error occurs or the end of the file is reached, the **getservbyname** subroutine returns a null pointer.

## Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

## Related Information

The **endprotoent** subroutine, **endservent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **getservbyport** subroutine, **getservent** subroutine, **setprotoent** subroutine, **setservent** subroutine.

Sockets Overview, and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## getservbyname\_r Subroutine

### Purpose

Gets service entry by name.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int getservbyname_r(name, proto, servent, serv_data)
const char *Name, proto;
struct servent *servent;
struct servent_data *serv_data;
```

### Description

**Requirement:** Use the **getservbyname** subroutine instead of the **getservbyname\_r** subroutine. The **getservbyname\_r** subroutine is compatible only with earlier versions of AIX.

An application program can use the **getservbyname\_r** subroutine to access a service, service aliases, the protocol for the service, and a protocol port number for the service.

The **getservbyname\_r** subroutine searches the **/etc/services** file sequentially from the start of the file until it finds one of the following:

- Matching name and protocol number.
- Matching name when the *Protocol* parameter is set to 0.
- End of the file.

Upon locating a matching name and protocol, the **getservbyname\_r** subroutine stores the values to the servent structure. The **getservbyname** subroutine overwrites the static data it returns in subsequent calls. The **getservbyname\_r** subroutine does not.

Use the **endservent\_r** subroutine to close the **/etc/hosts** file.

You must fill the **servent\_data** structure with zeros before its first access by either the **setservent\_r** or the **getservbyname\_r** subroutine.

## Parameters

<i>name</i>	Specifies the name of a service.
<i>proto</i>	Specifies a protocol for use with the specified service.
<i>servent</i>	Points to the <b>servent</b> structure.
<i>serv_data</i>	Points to the <b>serv_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful. The **getservbyname** subroutine returns a pointer to a servent structure when a successful match occurs. Entries in this structure are in network byte order.

**Note:** If an error occurs or the end of the file is reached, the **getservbyname\_r** returns a -1.

## Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

## Related Information

“endservent\_r Subroutine” on page 51, “setservent\_r Subroutine” on page 215, “getservent\_r Subroutine” on page 109, and “getservbyport\_r Subroutine” on page 106.

---

## getservbyport Subroutine

### Purpose

Gets service entry by port.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
struct servent *getservbyport (Port, Protocol)
int Port;char *Protocol;
```

### Description

The **getservbyport** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getservbyport** subroutine retrieves an entry from the **/etc/services** file using a port number as a search key.

An application program can use the **getservbyport** subroutine to access a service, service aliases, the protocol for the service, and a protocol port number for the service.

The **getservbyport** subroutine searches the services file sequentially from the beginning of the file until it finds one of the following:

- Matching protocol and port number
- Matching protocol when the *Port* parameter value equals 0
- End of the file

Upon locating a matching protocol and port number or upon locating a matching protocol only if the *Port* parameter value equals 0, the **getservbyport** subroutine returns a pointer to a **servent** structure, which contains fields for a line of information in the **/etc/services** file. The **netdb.h** file defines the **servent** structure and structure fields.

Use the **endservent** subroutine to close the **/etc/services** file.

All applications containing the **getservbyport** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Port</i>	Specifies the port where a service resides.
<i>Protocol</i>	Specifies a protocol for use with the service.

## Return Values

Upon successful completion, the **getservbyport** subroutine returns a pointer to a **servent** structure.

If an error occurs or the end of the file is reached, the **getservbyport** subroutine returns a null pointer.

## Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

## Related Information

The **endprotoent** subroutine, **endservent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **getservbyname** subroutine, **getservent** subroutine, **setprotoent** subroutine, **setservent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getservbyport\_r Subroutine

### Purpose

Gets service entry by port.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <netdb.h>
```

```
int getservbyport_r(Port, Proto, servent, serv_data)
int Port;
const char *Proto;
struct servent *servent;
struct servent_data *serv_data;
```

## Description

The **getservbyport\_r** subroutine retrieves an entry from the **/etc/services** file using a port number as a search key. An application program can use the **getservbyport\_r** subroutine to access a service, service aliases, the protocol for the service, and a protocol port number for the service.

The **getservbyport\_r** subroutine searches the services file sequentially from the beginning of the file until it finds one of the following:

- Matching protocol and port number
- Matching protocol when the *Port* parameter value equals 0
- End of the file

Upon locating a matching protocol and port number or upon locating a matching protocol where the *Port* parameter value equals 0, the **getservbyport\_r** subroutine returns a pointer to a servent structure, which contains fields for a line of information in the **/etc/services** file. The **netdb.h** file defines the servent structure, the servent\_data structure, and their fields.

The **getservbyport** routine overwrites static data returned on subsequent calls. The **getservbyport\_r** routine does not.

Use the **endservent\_r** subroutine to close the **/etc/services** file.

## Parameters

<i>Port</i>	Specifies the port where a service resides.
<i>Proto</i>	Specifies a protocol for use with the service.
<i>servent</i>	Points to the <b>servent</b> structure.
<i>serv_data</i>	Points to the <b>serv_data</b> structure.

## Return Values

The function returns a 0 if successful and a -1 if unsuccessful.

**Note:** If an error occurs or the end of the file is reached, the **getservbyport\_r** subroutine returns a -1 to indicate error.

## Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

## Related Information

“endservent\_r Subroutine” on page 51, “setservent\_r Subroutine” on page 215, “getservent\_r Subroutine” on page 109, and “getservbyname\_r Subroutine” on page 104.

---

## getservent Subroutine

### Purpose

Gets services file entry.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
struct servent *getservent ( )
```

### Description

The **getservent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **getservent** subroutine opens and reads the next line of the **/etc/services** file.

An application program can use the **getservent** subroutine to retrieve information about network services and the protocol ports they use.

The **getservent** subroutine returns a pointer to a **servent** structure, which contains fields for a line of information from the **/etc/services** file. The **servent** structure is defined in the **netdb.h** file.

The **/etc/services** file remains open after a call by the **getservent** subroutine. To close the **/etc/services** file after each call, use the **setservent** subroutine. Otherwise, use the **endservent** subroutine to close the **/etc/services** file.

All applications containing the **getservent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Return Values

The **getservent** subroutine returns a pointer to a **servent** structure when a successful match occurs.

If an error occurs or the end of the file is reached, the **getservent** subroutine returns a null pointer.

### Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

### Related Information

The **endprotoent** subroutine, **endservent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **setprotoent** subroutine, **setservent** subroutine.

Sockets Overview, and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## getservent\_r Subroutine

### Purpose

Gets services file entry.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int getservent_r(servent, serv_data)
struct servent *servent;
struct servent_data *serv_data;
```

### Description

The **getservent\_r** subroutine opens and reads the next line of the **/etc/services** file. An application program can use the **getservent\_r** subroutine to retrieve information about network services and the protocol ports they use.

The **/etc/services** file remains open after a call by the **getservent\_r** subroutine. To close the **/etc/services** file after each call, use the **setservent\_r** subroutine. Otherwise, use the **endservent\_r** subroutine to close the **/etc/services** file.

### Parameters

<i>servent</i>	Points to the <b>servent</b> structure.
<i>serv_data</i>	Points to the <b>serv_data</b> structure.

### Return Values

The **getservent\_r** fails when a successful match occurs. The **getservent** subroutine overwrites static data returned on subsequent calls. The **getservent\_r** subroutine does not.

### Files

<b>/etc/services</b>	Contains service names.
----------------------	-------------------------

### Related Information

“endservent\_r Subroutine” on page 51, “setservent\_r Subroutine” on page 215, “getservbyport\_r Subroutine” on page 106, and “getservbyname\_r Subroutine” on page 104.

---

## getsockname Subroutine

### Purpose

Gets the socket name.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int getsockname (Socket, Name, NameLength)
int Socket;
struct sockaddr * Name;
socklen_t * NameLength;
```

## Description

The **getsockname** subroutine retrieves the locally bound address of the specified socket. The socket address represents a port number in the Internet domain and is stored in the **sockaddr** structure pointed to by the *Name* parameter. The **sys/socket.h** file defines the **sockaddr** data structure.

**Note:** The **getsockname** subroutine does not perform operations on UNIX domain sockets.

A process created by another process can inherit open sockets. To use the inherited socket, the created process needs to identify their addresses. The **getsockname** subroutine allows a process to retrieve the local address bound to the specified socket.

A process can use the **getpeername** subroutine to determine the address of a destination socket in a socket connection.

All applications containing the **getsockname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Socket</i>	Specifies the socket for which the local address is desired.
<i>Name</i>	Points to the structure containing the local address of the specified socket.
<i>NameLength</i>	Specifies the size of the local address in bytes. Initializes the value pointed to by the <i>NameLength</i> parameter to indicate the amount of space pointed to by the <i>Name</i> parameter.

## Return Values

Upon successful completion, a value of 0 is returned, and the *NameLength* parameter points to the size of the socket address.

If the **getsockname** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **getsockname** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ENOBUFS</b>	Insufficient resources are available in the system to complete the call.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.

## Related Information

The **accept** subroutine, **bind** subroutine, **getpeername** subroutine, **socket** subroutine.

## getsockopt Subroutine

### Purpose

Gets options on sockets.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int getsockopt (Socket, Level, OptionName, OptionValue, OptionLength)
int Socket, Level, OptionName;
void * OptionValue;
socklen_t * OptionLength;
```

### Description

The **getsockopt** subroutine allows an application program to query socket options. The calling program specifies the name of the socket, the name of the option, and a place to store the requested information. The operating system gets the socket option information from its internal data structures and passes the requested information back to the calling program.

Options can exist at multiple protocol levels. They are always present at the uppermost socket level. When retrieving socket options, specify the level where the option resides and the name of the option.

All applications containing the **getsockopt** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the unique socket name.
<i>Level</i>	Specifies the protocol level where the option resides. Options can be retrieved at the following levels:

#### Socket level

Specifies the *Level* parameter as the **SOL\_SOCKET** option.

#### Other levels

Supplies the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP protocol, set the *Level* parameter to the protocol number of TCP, as defined in the **netinet/in.h** file.

*OptionName* Specifies a single option. The *OptionName* parameter and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The **sys/socket.h** file contains definitions for socket level options. The **netinet/tcp.h** file contains definitions for TCP protocol level options. Socket-level options can be enabled or disabled; they operate in a toggle fashion. The **sys/atmsock.h** file contains definitions for ATM protocol level options.

The following list defines socket protocol level options found in the **sys/socket.h** file:

#### **SO\_DEBUG**

Specifies the recording of debugging information. This option enables or disables debugging in the underlying protocol modules.

#### **SO\_BROADCAST**

Specifies whether transmission of broadcast messages is supported. The option enables or disables broadcast support.

#### **SO\_CKSUMREV**

Enables performance enhancements in the protocol layers. If the protocol supports this option, enabling causes the protocol to defer checksum verification until the user's data is moved into the user's buffer (on **recv**, **recvfrom**, **read**, or **recvmsg** thread). This can cause applications to be awakened when no data is available, in the case of a checksum error. In this case, **EAGAIN** is returned. Applications that set this option must handle the **EAGAIN** error code returned from a receive call.

#### **SO\_REUSEADDR**

Specifies that the rules used in validating addresses supplied by a **bind** subroutine should allow reuse of a local port. A particular IP address can only be bound once to the same port. This option enables or disables reuse of local ports.

**SO\_REUSEADDR** allows an application to explicitly deny subsequent **bind** subroutine to the port/address of the socket with **SO\_REUSEADDR** set. This allows an application to block other applications from binding with the **bind** subroutine.

#### **SO\_REUSEPORT**

Specifies that the rules used in validating addresses supplied by a **bind** subroutine should allow reuse of a local port/address combination. Each binding of the port/address combination must specify the **SO\_REUSEPORT** socket option. This option enables or disables the reuse of local port/address combinations.

#### **SO\_KEEPAIVE**

Monitors the activity of a connection by enabling or disabling the periodic transmission of ACK messages on a connected socket. The idle interval time can be designated using the TCP/IP **no** command. Broken connections are discussed in "Understanding Socket Types and Protocols" in *AIX Version 6.1 Communications Programming Concepts*.

#### **SO\_DONTROUTE**

Indicates outgoing messages should bypass the standard routing facilities. Does not apply routing on outgoing messages. Directs messages to the appropriate network interface according to the network portion of the destination address. This option enables or disables routing of outgoing messages.

#### **SO\_LINGER**

Lingers on a **close** subroutine if data is present. This option controls the action taken when an unsent messages queue exists for a socket, and a process performs a **close** subroutine on the socket.

If the **SO\_LINGER** option is set, the system blocks the process during the **close** subroutine until it can transmit the data or until the time expires. If the **SO\_LINGER** option is not specified, and a **close** subroutine is issued, the system handles the call in a way that allows the process to continue as quickly as possible.

The **sys/socket.h** file defines the **linger** structure that contains the **l\_linger** member for specifying linger time interval. If linger time is set to anything but 0, the system tries to send any messages queued on the socket. The maximum value that the **l\_linger** member can be set to is 65535. If the application has requested SPEC1170 compliant behavior by exporting the **XPG\_SUS\_ENV** environment variable, the linger time is *n* seconds; otherwise, the linger time is *n*/100 seconds (ticks), where *n* is the value of the **l\_linger** member.

#### **SO\_OOBINLINE**

Leaves received out-of-band data (data marked urgent) in line. This option enables or disables the receipt of out-of-band data.

#### **SO\_SNDBUF**

Retrieves buffer size information.

#### **SO\_RCVBUF**

Retrieves buffer size information.

#### **SO\_SNDLOWAT**

Retrieves send buffer low-water mark information.

#### **SO\_RCVLOWAT**

Retrieves receive buffer low-water mark information.

**SO\_SNDTIMEO**

Retrieves time-out information. This option is settable, but currently not used.

**SO\_RCVTIMEO**

Retrieves time-out information. This option is settable, but currently not used.

**SO\_PEERID**

Retrieves the credential information of the process associated with a peer UNIX domain socket. This information includes the process ID, effective user ID, and effective group ID. The **peercred\_struct** structure must be used in order to get the credential information. This structure is defined in the **sys/socket.h** file.

**SO\_ERROR**

Retrieves information about error status and clears.

The following list defines TCP protocol level options found in the **netinet/tcp.h** file:

**TCP\_RFC1323**

Indicates whether RFC 1323 is enabled or disabled on the specified socket. A non-zero *OptionValue* returned by the **getsockopt** subroutine indicates the RFC is enabled.

**TCP\_NODELAY**

Specifies whether TCP should follow the Nagle algorithm for deciding when to send data. By default TCP will follow the Nagle algorithm. To disable this behavior, applications can enable **TCP\_NODELAY** to force TCP to always send data immediately. A non-zero *OptionValue* returned by the **getsockopt** subroutine indicates **TCP\_NODELAY** is enabled. For example, **TCP\_NODELAY** should be used when there is an application using TCP for a request/response.

**TCP\_NODELAYACK**

Specifies if TCP needs to send immediate acknowledgement packets to the sender. If this option is not set, TCP delays sending the acknowledgement packets by up to 200 ms. This allows the acknowledgements to be sent along with the data on a response and minimizes system overhead. Setting this TCP option might cause a slight increase in system overhead, but can result in higher performance for network transfers if the sender is waiting on the receiver's acknowledgements.

The following list defines ATM protocol level options found in the **sys/atmsock.h** file:

**SO\_ATM\_PARM**

Retrieves all ATM parameters. This socket option can be used instead of using individual sockets options described below. It uses the **connect\_ie** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_AAL\_PARM**

Retrieves ATM AAL (Adaptation Layer) parameters. It uses the **aal\_parm** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_TRAFFIC\_DESC**

Retrieves ATM Traffic Descriptor values. It uses the **traffic\_desc** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BEARER**

Retrieves ATM Bearer capability information. It uses the **bearer** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BHLI**

Retrieves ATM Broadband High Layer Information. It uses the **bhli** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BLLI**

Retrieves ATM Broadband Low Layer Information. It uses the **blli** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_QoS**

Retrieves ATM Quality Of Service values. It uses the **qos\_parm** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_TRANSIT\_SEL**

Retrieves ATM Transit Selector Carrier. It uses the **transit\_sel** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_MAX\_PEND**

Retrieves the number of outstanding transmit buffers that are permitted before an error indication is returned to applications as a result of a transmit operation. This option is only valid for non best effort types of virtual circuits.

**SO\_ATM\_CAUSE**

Retrieves cause for the connection failure. It uses the **cause\_t** structure defined in the **sys/call\_ie.h** file.

*OptionValue* Specifies a pointer to the address of a buffer. The *OptionValue* parameter takes an integer parameter. The *OptionValue* parameter should be set to a nonzero value to enable a Boolean option or to a value of 0 to disable the option. The following options enable and disable in the same manner:

- **SO\_DEBUG**
- **SO\_REUSEADDR**
- **SO\_KEEPALIVE**
- **SO\_DONTROUTE**
- **SO\_BROADCAST**
- **SO\_OOBINLINE**
- **TCP\_RFC1323**

*OptionLength* Specifies the length of the *OptionValue* parameter. The *OptionLength* parameter initially contains the size of the buffer pointed to by the *OptionValue* parameter. On return, the *OptionLength* parameter is modified to indicate the actual size of the value returned. If no option value is supplied or returned, the *OptionValue* parameter can be 0.

Options at other protocol levels vary in format and name.

IP level (**IPPROTO\_IP** level) options are defined as follows:

**IP\_DONTFRAG**      Get current **IP\_DONTFRAG** option value.  
**IP\_FINDPMTU**      Get current PMTU value.  
**IP\_PMTUAGE**      Get current PMTU time out value.

In the case of TCP protocol sockets:

**IP\_DONTGRAG**      Not supported.  
**IP\_FINDPMTU**      Get current PMTU value.  
**IP\_PMTUAGE**      Not supported.

IPv6 level (**IPPROTO\_IPV6** level) options are defined as follows:

<b>IPV6_V6ONLY</b>	Determines whether the socket is restricted to IPv6 communications only.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_UNICAST_HOPS</b>	Allows the user to determine the outgoing hop limit value for unicast IPv6 packets.	
	<b>Option Type:</b>	int
<b>IPV6_MULTICAST_HOPS</b>	Allows the user to determine the outgoing hop limit value for multicast IPv6 packets.	
	<b>Option Type:</b>	int
<b>IPV6_MULTICAST_IF</b>	Allows the user to determine the interface being used for outgoing multicast packets.	
	<b>Option Type:</b>	unsigned int
<b>IPV6_MULTICAST_LOOP</b>	If a multicast datagram is sent to a group that the sending host belongs, a copy of the datagram is looped back by the IP layer for local delivery (if the option is set to 1). If the option is set to 0, a copy is not looped back.	
	<b>Option Type:</b>	unsigned int
<b>IPV6_RECVPKTINFO</b>	Determines whether the destination IPv6 address and arriving interface index of incoming IPv6 packets are being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_RECVHOPLIMIT</b>	Determines whether the hop limit of incoming IPv6 packets is being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_RECVTCLASS</b>	Determines whether the traffic class of incoming IPv6 packets is being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_RECVRTHDR</b>	Determines whether the routing header of incoming IPv6 packets is being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_RECVHOPOPTS</b>	Determine whether the hop-by-hop options header of incoming IPv6 packets is being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_RECVDSTOPTS</b>	Determines whether the destination options header of incoming IPv6 packets is being received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
<b>IPV6_PKTINFO</b>	Determines the source IPv6 address and outgoing interface index for all IPv6 packets being sent on this socket.	
	<b>Option Type:</b>	struct in6_pktinfo defined in the <b>netinet/in.h</b> file.

<b>IPV6_NEXTHOP</b>	Determines the next hop being used for outgoing IPv6 datagrams on this socket.	
	<b>Option Type:</b>	<b>struct sockaddr_in6</b> defined in the <b>netinet/in.h</b> file.
<b>IPV6_TCLASS</b>	Determines the traffic class for outgoing IPv6 datagrams on this socket.	
	<b>Option Type:</b>	<b>int</b>
<b>IPV6_RTHDR</b>	Determines the routing header to be used for outgoing IPv6 datagrams on this socket.	
	<b>Option Type:</b>	<b>struct ip6_rthdr</b> defined in the <b>netinet/ip6.h</b> file.
<b>IPV6_HOPOPTS</b>	Determines the hop-by-hop options header to be used for outgoing IPv6 datagrams on this socket.	
	<b>Option Type:</b>	<b>struct ip6_hbh</b> defined in the <b>netinet/ip6.h</b> file.
<b>IPV6_DSTOPTS</b>	Determines the destination options header to be used for outgoing IPv6 datagrams on this socket. This header will follow a routing header (if present) and will also be used when there is no routing header specified.	
	<b>Option Type:</b>	<b>struct ip6_dest</b> defined in the <b>netinet/ip6.h</b> file.
<b>IPV6_RTHDRDSTOPTS</b>	Determines the destination options header to be used for outgoing IPv6 datagrams on this socket. This header will precede a routing header (if present). If no routing header is specified, this option will be silently ignored.	
	<b>Option Type:</b>	<b>struct ip6_dest</b> defined in the <b>netinet/ip6.h</b> file.
<b>IPV6_USE_MIN_MTU</b>	Determines how IPv6 path MTU discovery is being controlled for this socket.	
	<b>Option Type:</b>	<b>int</b>
<b>IPV6_DONTFRAG</b>	Determines whether fragmentation of outgoing IPv6 packets has been disabled on this socket.	
	<b>Option Type:</b>	<b>int</b> (boolean interpretation)
<b>IPV6_RECVPATHMTU</b>	Determines whether IPV6_PATHMTU messages are being received as ancillary data on this socket.	
	<b>Option Type:</b>	<b>int</b> (boolean interpretation)
<b>IPV6_PATHMTU</b>	Determines the current Path MTU for a connected socket.	
	<b>Option Type:</b>	<b>struct ip6_mtuinfo</b> defined in the <b>netinet/in.h</b> file.

ICMPV6 level (**IPPROTO\_ICMPV6** level) options are defined as follows:

<b>ICMP6_FILTER</b>	Allows the user to filter ICMPV6 messages by the ICMPV6 type field. If no filter was set, the default kernel filter will be returned.	
	<b>Option Type:</b>	The <b>icmp6_filter</b> structure defined in the <b>netinet/icmp6.h</b> file.

## Return Values

Upon successful completion, the **getsockopt** subroutine returns a value of 0.

If the **getsockopt** subroutine is unsuccessful, the subroutine handler performs the following actions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

Upon successful completion of the **IPPROTO\_IP** option **IP\_PMTUAGE** the returns are:

Prior to AIX 5.3:

- OptionValue 0 if PMTU discovery is not enabled.
- OptionValue -1 if PMTU discovery is not complete.

- Positive non-zero OptionValue if PMTU is available.

Beginning with AIX 5.3:

- Positive non-zero OptionValue.

Upon successful completion of TCP protocol sockets option **IP\_FINDPMTU** the returns are:

Prior to AIX 5.3:

- OptionValue 0 if PMTU discovery (tcp\_pmtu\_discover) is not enabled.
- OptionValue -1 if PMTU discovery is not complete/not available.
- Positive non-zero OptionValue if PMTU is available.

Beginning with AIX 5.3:

- OptionValue 0 if PMTU discovery (tcp\_pmtu\_discover) is not enabled/not available.
- Positive non-zero OptionValue if PMTU is available.

## Error Codes

The **getsockopt** subroutine is unsuccessful if any of the following errors occurs:

<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EFAULT</b>	The address pointed to by the <i>OptionValue</i> parameter is not in a valid (writable) part of the process space, or the <i>OptionLength</i> parameter is not in a valid part of the process address space.
<b>EINVAL</b>	The <i>Level</i> , <i>OptionName</i> , or <i>OptionLength</i> is invalid.
<b>ENOBUF</b>	Insufficient resources are available in the system to complete the call.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ENOPROTOOPT</b>	The option is unknown.
<b>EOPNOTSUPP</b>	The option is not supported by the socket family or socket type.

## Examples

The following program fragment illustrates the use of the **getsockopt** subroutine to determine an existing socket type:

```
#include <sys/types.h>
#include <sys/socket.h>
int type, size;
size = sizeof(int);
if(getsockopt(s, SOL_SOCKET, SO_TYPE, (char*)&type,&size)<0){
.
.
.
}
```

## Related Information

The **no** command.

The **bind** subroutine, **close** subroutine, **endprotoent** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **setprotoent** subroutine, **setsockopt** subroutine, **socket** subroutine.

Sockets Overview, Understanding Socket Options, and Understanding Socket Types and Protocols in *AIX Version 6.1 Communications Programming Concepts*.

---

## htonl Subroutine

### Purpose

Converts an unsigned long integer from host byte order to Internet network byte order.

### Library

ISODE Library (**libisode.a**)

### Syntax

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

```
uint32_t htonl ( HostLong)
uint32_t HostLong;
```

### Description

The **htonl** subroutine converts an unsigned long (32-bit) integer from host byte order to Internet network byte order.

The Internet network requires addresses and ports in network standard byte order. Use the **htonl** subroutine to convert the host integer representation of addresses and ports to Internet network byte order.

The **htonl** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is the same as the network byte order.

The **htonl** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not the same as the network byte order.

All applications containing the **htonl** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

*HostLong*            Specifies a 32-bit integer in host byte order.

### Return Values

The **htonl** subroutine returns a 32-bit integer in Internet network byte order (most significant byte first).

### Related Information

The **htons** subroutine, **ntohl** subroutine, **ntohs** subroutine, **htonll** subroutine, **ntohll** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## htonll Subroutine

### Purpose

Converts an unsigned long integer from host byte order to Internet network byte order.

## Library

ISODE Library (**libisode.a**)

## Syntax

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

```
uint64_t htonl ( HostLong)
uint64_t HostLong;
```

## Description

The **htonll** subroutine converts an unsigned long (64-bit) integer from host byte order to Internet network byte order.

The Internet network requires addresses and ports in network standard byte order. Use the **htonll** subroutine to convert the host integer representation of addresses and ports to Internet network byte order.

The **htonll** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is the same as the network byte order.

The **htonll** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not the same as the network byte order.

All applications containing the **htonll** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*HostLong*            Specifies a 64-bit integer in host byte order.

## Return Values

The **htonll** subroutine returns a 64-bit integer in Internet network byte order (most significant byte first).

## Related Information

The **htons** subroutine, **ntohl** subroutine, **ntohs** subroutine, **htonll** subroutine, **ntohl** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## htons Subroutine

### Purpose

Converts an unsigned short integer from host byte order to Internet network byte order.

## Library

ISODE Library (**libisode.a**)

## Syntax

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

```
uint16_t htons ( HostShort)
uint16_t HostShort;
```

## Description

The **htons** subroutine converts an unsigned short (16-bit) integer from host byte order to Internet network byte order.

The Internet network requires ports and addresses in network standard byte order. Use the **htons** subroutine to convert addresses and ports from their host integer representation to network standard byte order.

The **htons** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is the same as the network byte order.

The **htons** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not the same as the network byte order.

All applications containing the **htons** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*HostShort*                Specifies a 16-bit integer in host byte order that is a host address or port.

## Return Values

The **htons** subroutine returns a 16-bit integer in Internet network byte order (most significant byte first).

## Related Information

The **htonl** subroutine, **ntohl** subroutine, **ntohs** subroutine, **htonll** subroutine, **ntohll** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## if\_freenameindex Subroutine

### Purpose

Frees the dynamic memory that was allocated by the “if\_nameindex Subroutine” on page 123.

### Library

Library (**libc.a**)

### Syntax

```
#include <net/if.h>
```

```
void if_freenameindex (struct if_nameindex *ptr);
```

### Description

The *ptr* parameter is a pointer returned by the **if\_nameindex** subroutine. After the **if\_freenameindex** subroutine has been called, the application must not use the array of which *ptr* is the address.

## Parameters

*ptr*

Pointer returned by the **if\_nameindex** subroutine

## Related Information

“if\_nametoindex Subroutine” on page 123, “if\_indextoname Subroutine,” and “if\_nameindex Subroutine” on page 123.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## if\_indextoname Subroutine

### Purpose

Maps an interface index into its corresponding name.

### Library

Standard C Library <**libc.a**>

### Syntax

```
#include <net/if.h>
char *if_indextoname(unsigned int ifindex, char *ifname);
```

### Description

When the **if\_indextoname** subroutine is called, the *ifname* parameter points to a buffer of at least IF\_NAMESIZE bytes. The **if\_indextoname** subroutine places the name of the interface in this buffer with the *ifindex* index.

**Note:** IF\_NAMESIZE is also defined in <**net/if.h**> and its value includes a terminating null byte at the end of the interface name.

If *ifindex* is an interface index, the **if\_indextoname** Subroutine returns the *ifname* value, which points to a buffer containing the interface name. Otherwise, it returns a NULL pointer and sets the **errno** global value to indicate the error.

If there is no interface corresponding to the specified index, the **errno** global value is set to **ENXIO**. If a system error occurs (such as insufficient memory), the **errno** global value is set to the proper value (such as, **ENOMEM**).

## Parameters

*ifindex*

Possible interface index

*ifname*

Possible name of an interface

## Error Codes

**ENXIO**

There is no interface corresponding to the specified index

**ENOMEM**

Insufficient memory

## Related Information

“if\_nametoindex Subroutine,” “if\_indextoname Subroutine” on page 122, and “if\_nameindex Subroutine.”

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## if\_nameindex Subroutine

### Purpose

Retrieves index and name information for all interfaces.

### Library

The Standard C Library (<libc.a>)

### Syntax

```
#include <net/if.h>

struct if_nameindex *if_nameindex(void)
struct if_nameindex {
    unsigned int if_index; /* 1, 2, ... */
    char *if_name; /* null terminated name: "1e0", ... */
};
```

### Description

The **if\_nameindex** subroutine returns an array of **if\_nameindex** structures (one per interface).

The memory used for this array of structures is obtained dynamically. The interface names pointed to by the *if\_name* members are obtained dynamically as well. This memory is freed by the **if\_freenameindex** subroutine.

The function returns a NULL pointer upon error, and sets the **errno** global value to the appropriate value. If successful, the function returns an array of structures. The end of an array of structures is indicated by a structure with an *if\_index* value of 0 and an *if\_name* value of NULL.

## Related Information

“if\_nametoindex Subroutine,” “if\_indextoname Subroutine” on page 122, and “if\_freenameindex Subroutine” on page 121.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## if\_nametoindex Subroutine

### Purpose

Maps an interface name to its corresponding index.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <net/if.h>
unsigned int if_nametoindex(const char *ifname);
```

## Description

If the *ifname* parameter is the name of an interface, the **if\_nametoindex** subroutine returns the interface index corresponding to the *ifname* name. If the *ifname* parameter is not the name of an interface, the **if\_nametoindex** subroutine returns a 0 and the **errno** global variable is set to the appropriate value.

## Parameters

*ifname* Possible name of an interface.

## Related Information

“if\_indextoname Subroutine” on page 122, “if\_nameindex Subroutine” on page 123, and “if\_freenameindex Subroutine” on page 121.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## inet6\_opt\_append Subroutine

### Purpose

Returns the updated total length of the extension header.

### Syntax

```
int inet6_opt_append(void *extbuf, socklen_t extlen, int offset,
                    uint8_t type, socklen_t len, uint_t align,
                    void **databufp);
```

### Description

The **inet6\_opt\_append** subroutine returns the updated total length of the extension header, taking into account adding an option with length *len* and alignment *align*. If *extbuf* is not NULL, then, in addition to returning the length, the subroutine inserts any needed pad option, initializes the option (setting the type and length fields), and returns a pointer to the location for the option content in *databufp*. After **inet6\_opt\_append()** has been called, the application can use the *databuf* directly, or use **inet6\_opt\_set\_val()** to specify the content of the option.

### Parameters

<i>extbuf</i>	If NULL, <b>inet6_opt_append</b> will return only the updated length. If <i>extbuf</i> is not NULL, in addition to returning the length, the function inserts any needed pad option, initializes the option (setting the <b>type</b> and <b>length</b> fields) and returns a pointer to the location for the option content in <i>databufp</i> .
<i>extlen</i>	Size of the buffer pointed to by <i>extbuf</i> .
<i>offset</i>	The length returned by <b>inet6_opt_init()</b> or a previous <b>inet6_opt_append()</b> .
<i>type</i>	8-bit option type. Must have a value from 2 to 255, inclusive. (0 and 1 are reserved for the <i>Pad1</i> and <i>PadN</i> options, respectively.)

<i>len</i>	Length of the option data (excluding the option type and option length fields). Must be a value between 0 and 255, inclusive, and is the length of the option data that follows.
<i>align</i>	Alignment of the option data. Must be a value of 1, 2, 4, or 8. The <i>align</i> value can not exceed the value of <i>len</i> .
<i>databufp</i>	Specifies the content of the option.

## Return Values

-1	Option content does not fit in the extension header buffer.
integer value	Updated total length of the extension header.

## Related Information

“inet6\_opt\_find Subroutine,” “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_find Subroutine

### Purpose

Looks for a specified option in the extension header.

### Syntax

```
int inet6_opt_find(void *extbuf, socklen_t extlen, int offset,
                  uint8_t *typep, socklen_t *lenp,
                  void **databufp);
```

### Description

The **inet6\_opt\_find** subroutine is similar to the **inet6\_opt\_next()** function, except this subroutine lets the caller specify the option type to be searched for, instead of always returning the next option in the extension header.

### Parameters

<i>extbuf</i>	Specifies the extension header.
<i>extlen</i>	Size of the buffer pointed to by <i>extbuf</i> .
<i>offset</i>	Specifies the position where scanning of the extension buffer can continue. Should either be 0 (for the first option) or the length returned by a previous call to <b>inet6_opt_next()</b> or <b>inet6_opt_find()</b> .
<i>typep</i>	Stores the option type.
<i>lenp</i>	Stores the length of the option data (excluding the option type and option length fields).
<i>databufp</i>	Points to the data field of the option.

## Return Values

The **inet6\_opt\_find** subroutine returns the updated “previous” total length computed by advancing past the option that was returned and past any options that did not match the type. This returned “previous”

length can then be passed to subsequent calls to **inet6\_opt\_find()** for finding the next occurrence of the same option type.

-1                   The option cannot be located, there are no more options, or the option extension header is malformed.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_finish Subroutine,” “inet6\_opt\_get\_val Subroutine,” “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_finish Subroutine

### Purpose

Returns the final length of an extension header.

### Syntax

```
int inet6_opt_finish(void *extbuf, socklen_t extlen, int offset);
```

### Description

The **inet6\_opt\_finish** subroutine returns the final length of an extension header, taking into account the final padding of the extension header to make it a multiple of 8 bytes.

### Parameters

<i>extbuf</i>	If NULL, <b>inet6_opt_finish</b> will only return the final length. If <i>extbuf</i> is not NULL, in addition to returning the length, the function initializes the option by inserting a <i>Pad1</i> or <i>PadN</i> option of the proper length.
<i>extlen</i>	Size of the buffer pointed to by <i>extbuf</i> .
<i>offset</i>	The length returned by <b>inet6_opt_init()</b> or a previous <b>inet6_opt_append()</b> .

### Return Values

-1	The necessary pad does not fit in the extension header buffer.
integer value	Final length of the extension header.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_get\_val Subroutine,” “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_get\_val Subroutine

### Purpose

Extracts data items of various sizes in the data portion of the option.

## Syntax

```
int inet6_opt_get_val(void *databuf, int offset, void *val,
                     socklen_t vallen);
```

## Description

The **inet6\_opt\_get\_val** subroutine extracts data items of various sizes in the data portion of the option. It is expected that each field is aligned on its natural boundaries, but the subroutine will not rely on the alignment.

## Parameters

<i>databuf</i>	Pointer to the data content returned by <b>inet6_opt_next()</b> or <b>inet6_opt_find()</b> .
<i>offset</i>	Specifies where in the data portion of the option the value should be extracted. The first byte after the option type and length is accessed by specifying an offset of 0.
<i>val</i>	Pointer to the destination for the extracted data.
<i>vallen</i>	Specifies the size of the data content to be extracted.

## Return Values

The **inet6\_opt\_get\_val** subroutine returns the offset for the next field (that is, *offset* + *vallen*), which can be used when extracting option content with multiple fields.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_init Subroutine,” “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_init Subroutine

### Purpose

Returns the number of bytes needed for an empty extension header.

### Syntax

```
int inet6_opt_init(void *extbuf, socklen_t extlen);
```

### Description

The **inet6\_opt\_init** subroutine returns the number of bytes needed for the empty extension header (that is, a header without any options).

### Parameters

<i>extbuf</i>	Specifies NULL for an empty header. If <i>extbuf</i> is not NULL, it initializes the extension header to have the correct length field.
<i>extlen</i>	Specifies the size of the extension header. The value of <i>extlen</i> must be a positive value that is a multiple of 8.

### Return Values

-1                      The value of *extlen* is not a positive (non-zero) multiple of 8.

integer value

Number of bytes needed for an empty extension header.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_next Subroutine,” “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_next Subroutine

### Purpose

Parses received option extension headers returning the next option.

### Syntax

```
int inet6_opt_next(void *extbuf, socklen_t extlen, int offset,  
                  uint8_t *typep, socklen_t *lenp,  
                  void **databufp);
```

### Description

The **inet6\_opt\_next** subroutine parses received option extension headers, returning the next option. The next option is returned by updating the *typep*, *lenp*, and *databufp* parameters.

### Parameters

*extbuf*  
*extlen*  
*offset*

Specifies the extension header.  
Size of the buffer pointed to by *extbuf*.  
Specifies the position where scanning of the extension buffer can continue. Should either be 0 (for the first option) or the length returned by a previous call to **inet6\_opt\_next()** or **inet6\_opt\_find()**.

*typep*  
*lenp*

Stores the option type.  
Stores the length of the option data (excluding the option type and option length fields).

*databufp*

Points to the data field of the option.

### Return Values

The **inet6\_opt\_next** subroutine returns the updated “previous” length computed by advancing past the option that was returned. This returned “previous” length can then be passed to subsequent calls to **inet6\_opt\_next()**. This function does not return any *PAD1* or *PADN* options.

–1            There are no more options or the option extension header is malformed.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_opt\_set\_val Subroutine

### Purpose

Inserts data items into the data portion of an option.

### Syntax

```
int inet6_opt_set_val(void *databuf, int offset, void *val,
                     socklen_t vallen);
```

### Description

The **inet6\_opt\_set\_val** subroutine inserts data items of various sizes into the data portion of the option. The caller must ensure that each field is aligned on its natural boundaries. However, even when the alignment requirement is not satisfied, **inet6\_opt\_set\_val** will just copy the data as required.

### Parameters

<i>databuf</i>	Pointer to the data area returned by <b>inet6_opt_append()</b> .
<i>offset</i>	Specifies where in the data portion of the option the value should be inserted; the first byte after the option type and length is accessed by specifying an offset of 0.
<i>val</i>	Pointer to the data content to be inserted.
<i>vallen</i>	Specifies the size of the data content to be inserted.

### Return Values

The function returns the offset for the next field (that is, *offset* + *vallen*), which can be used when composing option content with multiple fields.

### Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_rth\_add Subroutine,” “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_add Subroutine

### Purpose

Adds an IPv6 address to the end of the Routing header being constructed.

### Syntax

```
int inet6_rth_add(void *bp, const struct in6_addr *addr);
```

### Description

The **inet6\_rth\_add** subroutine adds the IPv6 address pointed to by *addr* to the end of the Routing header being constructed.

### Parameters

<i>bp</i>	Points to the buffer of the Routing header.
<i>addr</i>	Specifies which IPv6 address is to be added.

## Return Values

0	Success. The <i>segleft</i> member of the Routing Header is updated to account for the new address in the Routing header.
-1	The new address could not be added.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine,” “inet6\_rth\_init Subroutine,” “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_getaddr Subroutine

### Purpose

Returns a pointer to a specific IPv6 address in a Routing header.

### Syntax

```
struct in6_addr *inet6_rth_getaddr(const void *bp, int index);
```

### Description

The **inet6\_rth\_getaddr** subroutine returns a pointer to the IPv6 address specified by *index* in the Routing header described by *bp*. An application should first call **inet6\_rth\_segments()** to obtain the number of segments in the Routing header.

### Parameters

<i>bp</i>	Points to the Routing header.
<i>index</i>	Specifies the index of the IPv6 address that must be returned. The value of <i>index</i> must be between 0 and one less than the value returned by <b>inet6_rth_segments()</b> .

## Return Values

NULL	The <b>inet6_rth_getaddr</b> subroutine failed.
Valid pointer	Pointer to the address indexed by <i>index</i> .

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_init Subroutine,” “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_init Subroutine

### Purpose

Initializes a buffer to contain a Routing header.

## Syntax

```
void *inet6_rth_init(void *bp, socklen_t bp_len, int type,
                    int segments);
```

## Description

The **inet6\_rth\_init** subroutine initializes the buffer pointed to by *bp* to contain a Routing header of the specified *type* and sets **ip6r\_len** based on the *segments* parameter. *bp\_len* is only used to verify that the buffer is large enough. The **ip6r\_segleft** field is set to 0; **inet6\_rth\_add()** increments it.

When the application uses ancillary data, the application must initialize any **cmsghdr** fields. The caller must allocate the buffer, and the size of the buffer can be determined by calling **inet6\_rth\_space()**.

## Parameters

<i>bp</i>	Points to the buffer to be initialized.
<i>bp_len</i>	Size of the buffer pointed to by <i>bp</i> .
<i>type</i>	Specifies the type of Routing header to be held.
<i>segments</i>	Specifies the number of addresses within the Routing header.

## Return Values

Upon success, the return value is the pointer to the buffer (*bp*), and this is then used as the first argument to the **inet6\_rth\_add()** function.

NULL                    The buffer could not be initialized.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_reverse Subroutine,” “inet6\_rth\_segments Subroutine” on page 132, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_reverse Subroutine

### Purpose

Writes a new Routing header that sends datagrams along the reverse route of a Routing header extension header.

### Syntax

```
int inet6_rth_reverse(const void *in, void *out);
```

### Description

The **inet6\_rth\_reverse** subroutine takes a Routing header extension header (pointed to by the first argument) and writes a new Routing header that sends datagrams along the reverse of that route. The function reverses the order of the addresses and sets the *segleft* member in the new Routing header to the number of segments. Both arguments are allowed to point to the same buffer (that is, the reversal can occur in place).

## Parameters

<i>in</i>	Points to the original Routing header extension header.
<i>out</i>	Points to the new Routing header route that reverses the route of <i>in</i> .

## Return Values

0	The reverse Routing header was successfully created.
-1	The reverse Routing header could not be created.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_segments Subroutine,” “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_segments Subroutine

### Purpose

Returns the number of segments (addresses) contained in a Routing header.

### Syntax

```
int inet6_rth_segments(const void *bp);
```

### Description

The **inet6\_rth\_segments** subroutine returns the number of segments (addresses) contained in the Routing header described by *bp*.

## Parameters

<i>bp</i>	Points to the Routing header.
-----------	-------------------------------

## Return Values

0 (or greater)	The number of addresses in the Routing header was returned.
-1	The number of addresses of the Routing header could not be returned.

## Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_space Subroutine” on page 133

---

## inet6\_rth\_space Subroutine

### Purpose

Returns the required number of bytes to hold a Routing header.

### Syntax

```
socklen_t inet6_rth_space(int type, int segments);
```

### Description

The **inet6\_rth\_space** subroutine returns the number of bytes required to hold a Routing header of the specified *type* containing the specified number of *segments* (addresses). For an IPv6 Type 0 Routing header, the number of segments must be between 0 and 127, inclusive. The return value is simply the space for the Routing header. When the application uses ancillary data, the application must pass the returned length to **CMSG\_SPACE()** in order to determine how much memory is needed for the ancillary data object (including the **cmsg\_hdr** structure).

**Note:** Although **inet6\_rth\_space** returns the size of the ancillary data, it does not allocate the space required for the ancillary data. This allows an application to allocate a larger buffer, so that other ancillary data objects can be added, because all the ancillary data objects must be specified to **sendmsg()** as a single **msg\_control** buffer.

### Parameters

<i>type</i>	Specifies the type of Routing header to be held.
<i>segments</i>	Specifies the number of addresses within the Routing header.

### Return Values

0	Either the type of the Routing header is not supported by this implementation or the number of segments is invalid for this type of Routing header.
<b>length</b>	Determines how much memory is needed for the ancillary data object.

### Related Information

“inet6\_opt\_append Subroutine” on page 124, “inet6\_opt\_find Subroutine” on page 125, “inet6\_opt\_finish Subroutine” on page 126, “inet6\_opt\_get\_val Subroutine” on page 126, “inet6\_opt\_init Subroutine” on page 127, “inet6\_opt\_next Subroutine” on page 128, “inet6\_opt\_set\_val Subroutine” on page 129, “inet6\_rth\_add Subroutine” on page 129, “inet6\_rth\_getaddr Subroutine” on page 130, “inet6\_rth\_init Subroutine” on page 130, “inet6\_rth\_reverse Subroutine” on page 131, “inet6\_rth\_segments Subroutine” on page 132

---

## inet\_addr Subroutine

### Purpose

Converts Internet addresses to Internet numbers.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <sys/socket.h>
#include <sys/socketvar.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
in_addr_t inet_addr ( CharString)
register const char *CharString;
```

## Description

The **inet\_addr** subroutine converts an ASCII string containing a valid Internet address using dot notation into an Internet address number typed as an unsigned integer value. An example of dot notation is 120.121.5.123. The **inet\_addr** subroutine returns an error value if the Internet address notation in the ASCII string supplied by the application is not valid.

**Note:** Although they both convert Internet addresses in dot notation to Internet numbers, the **inet\_addr** subroutine and **inet\_network** process ASCII strings differently. When an application gives the **inet\_addr** subroutine a string containing an Internet address value without a delimiter, the subroutine returns the logical product of the value represented by the string and 0xFFFFFFFF. For any other Internet address, if the value of the fields exceeds the previously defined limits, the **inet\_addr** subroutine returns an error value of -1.

When an application gives the **inet\_network** subroutine a string containing an Internet address value without a delimiter, the **inet\_network** subroutine returns the logical product of the value represented by the string and 0xFF. For any other Internet address, the subroutine returns an error value of -1 if the value of the fields exceeds the previously defined limits.

All applications containing the **inet\_addr** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

Sample return values for each subroutine are as follows:

Application String	inet_addr Returns	inet_network Returns
0x1234567890abcdef 0x1234567890abcdef. 256.257.258.259	0x090abcdef 0xFFFFFFFF (= -1) 0xFFFFFFFF (= -1)	0x000000ef 0x0000ef00 0x00010203

The ASCII string for the **inet\_addr** subroutine must conform to the following format:

```
string ::= field | field delimited_field^1-3 | delimited_field^1-3
delimited_field ::= delimiter field | delimiter
delimiter ::= .
field ::= 0X | 0x | 0Xhexadecimal* | 0x hexadecimal* | decimal* | 0 octal
hexadecimal ::= decimal | a|b|c|d|e|f|A|B|C|D|E|F
decimal ::= octal | 8|9
octal ::= 0|1|2|3|4|5|6|7
```

### Notes:

1. ^n indicates *n* repetitions of a pattern.
2. ^n-m indicates *n* to *m* repetitions of a pattern.
3. \* indicates 0 or more repetitions of a pattern, up to environmental limits.
4. The Backus Naur form (BNF) description states the space character, if one is used. *Text* indicates text, not a BNF symbol.

The **inet\_addr** subroutine requires an application to terminate the string with a null terminator (0x00) or a space (0x30). The string is considered invalid if the application does not end it with a null terminator or a space. The subroutine ignores characters trailing a space.

The following describes the restrictions on the field values for the **inet\_addr** subroutine:

Format	Field Restrictions (in decimal)
a	<i>Value_a</i> < 4,294,967,296
a.b	<i>Value_a</i> < 256; <i>Value_b</i> < 16,777,216
a.b.c	<i>Value_a</i> < 256; <i>Value_b</i> < 256; <i>Value_c</i> < 65536
a.b.c.d	<i>Value_a</i> < 256; <i>Value_b</i> < 256; <i>Value_c</i> < 256; <i>Value_d</i> < 256

Applications that use the **inet\_addr** subroutine can enter field values exceeding these restrictions. The subroutine accepts the least significant bits up to an integer in length, then checks whether the truncated value exceeds the maximum field value. For example, if an application enters a field value of 0x1234567890 and the system uses 16 bits per integer, then the **inet\_addr** subroutine uses bits 0 -15. The subroutine returns 0x34567890.

Applications can omit field values between delimiters. The **inet\_addr** subroutine interprets empty fields as 0.

#### Notes:

1. The **inet\_addr** subroutine does not check the pointer to the ASCII string. The user must ensure the validity of the address in the ASCII string.
2. The application must verify that the network and host IDs for the Internet address conform to either a Class A, B, or C Internet address. The **inet\_atrr** subroutine processes any other number as a Class C address.

## Parameters

*CharString* Represents a string of characters in the Internet address form.

## Return Values

For valid input strings, the **inet\_addr** subroutine returns an unsigned integer value comprised of the bit patterns of the input fields concatenated together. The subroutine places the first pattern in the most significant position and appends any subsequent patterns to the next most significant positions.

The **inet\_addr** subroutine returns an error value of -1 for invalid strings.

**Note:** An Internet address with a dot notation value of 255.255.255.255 or its equivalent in a different base format causes the **inet\_addr** subroutine to return an unsigned integer value of 4294967295. This value is identical to the unsigned representation of the error value. Otherwise, the **inet\_addr** subroutine considers 255.255.255.255 a valid Internet address.

## Files

**/etc/hosts**

Contains host names.

**/etc/networks**

Contains network names.

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_Inaof** subroutine, **inet\_makeaddr** subroutine, **inet\_netof** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_Inaof Subroutine

### Purpose

Returns the host ID of an Internet address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
int inet_Inaof ( InternetAddr)
struct in_addr InternetAddr;
```

### Description

The **inet\_Inaof** subroutine masks off the host ID of an Internet address based on the Internet address class. The calling application must enter the Internet address as an unsigned long value.

All applications containing the **inet\_Inaof** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

**Note:** The application must verify that the network and host IDs for the Internet address conform to either a Class A, B, or C Internet address. The **inet\_Inaof** subroutine processes any other number as a Class C address.

### Parameters

*InternetAddr*                      Specifies the Internet address to separate.

### Return Values

The return values of the **inet\_Inaof** subroutine depend on the class of Internet address the application provides:

Value	Description
Class A	The logical product of the Internet address and 0x00FFFFFF.
Class B	The logical product of the Internet address and 0x0000FFFF.
Class C	The logical product of the Internet address and 0x000000FF.

## Files

**/etc/hosts**

Contains host names.

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_makeaddr** subroutine, **inet\_netof** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine. **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_makeaddr Subroutine

### Purpose

Returns a structure containing an Internet address based on a network ID and host ID provided by the application.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
struct in_addr inet_makeaddr ( Net, LocalNetAddr)
int Net, LocalNetAddr;
```

### Description

The **inet\_makeaddr** subroutine forms an Internet address from the network ID and Host ID provided by the application (as integer types). If the application provides a Class A network ID, the **inet\_makeaddr** subroutine forms the Internet address using the net ID in the highest-order byte and the logical product of the host ID and 0x00FFFFFF in the 3 lowest-order bytes. If the application provides a Class B network ID, the **inet\_makeaddr** subroutine forms the Internet address using the net ID in the two highest-order bytes and the logical product of the host ID and 0x0000FFFF in the lowest two ordered bytes. If the application does not provide either a Class A or Class B network ID, the **inet\_makeaddr** subroutine forms the Internet address using the network ID in the 3 highest-order bytes and the logical product of the host ID and 0x0000FFFF in the lowest-ordered byte.

The **inet\_makeaddr** subroutine ensures that the Internet address format conforms to network order, with the first byte representing the high-order byte. The **inet\_makeaddr** subroutine stores the Internet address in the structure as an unsigned long value.

The application must verify that the network ID and host ID for the Internet address conform to class A, B, or C. The **inet\_makeaddr** subroutine processes any nonconforming number as a Class C address.

The **inet\_makeaddr** subroutine expects the **in\_addr** structure to contain only the Internet address field. If the application defines the **in\_addr** structure otherwise, then the value returned in **in\_addr** by the **inet\_makeaddr** subroutine is undefined.

All applications containing the **inet\_makeaddr** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Net</i>	Contains an Internet network number.
<i>LocalNetAddr</i>	Contains a local network address.

## Return Values

Upon successful completion, the **inet\_makeaddr** subroutine returns a structure containing an Internet address.

If the **inet\_makeaddr** subroutine is unsuccessful, the subroutine returns a -1.

## Files

<i>/etc/hosts</i>	Contains host names.
-------------------	----------------------

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_netof** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_net\_ntop Subroutine

### Purpose

Converts between binary and text address formats.

### Library

Library (**libc.a**)

### Syntax

```
char *inet_net_ntop (af, src, bits, dst, size)
int af;
const void *src;
int bits;
char *dst;
size_t size;
```

## Description

This function converts a network address and the number of bits in the network part of the address into the CIDR format ascii text (for example, 9.3.149.0/24). The *af* parameter specifies the family of the address. The *src* parameter points to a buffer holding an IPv4 address if the *af* parameter is AF\_INET. The *bits* parameter is the size (in bits) of the buffer pointed to by the *src* parameter. The *dst* parameter points to a buffer where the function stores the resulting text string. The *size* parameter is the size (in bytes) of the buffer pointed to by the *dst* parameter.

## Parameters

<i>af</i>	Specifies the family of the address.
<i>src</i>	Points to a buffer holding an IPv4 address if the <i>af</i> parameter is AF_INET.
<i>bits</i>	Specifies the size of the buffer pointed to by the <i>src</i> parameter.
<i>dst</i>	Points to a buffer where the resulting text string is stored.
<i>size</i>	Specifies the size of the buffer pointed to by the <i>dst</i> parameter.

## Return Values

If successful, a pointer to a buffer containing the text string is returned. If unsuccessful, NULL is returned. Upon failure, **errno** is set to EAFNOSUPPORT if the *af* parameter is invalid or ENOSPC if the size of the result buffer is inadequate.

## Related Information

“inet\_net\_pton Subroutine,” “inet\_ntop Subroutine” on page 144, and “inet\_pton Subroutine” on page 145.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## inet\_net\_pton Subroutine

### Purpose

Converts between text and binary address formats.

### Library

Library (**libc.a**)

### Syntax

```
int inet_net_pton (af, src, dst, size)
int af;
const char *src;
void *dst;
size_t size;
```

### Description

This function converts a network address in ascii into the binary network address. The ascii representation can be CIDR-based (for example, 9.3.149.0/24) or class-based (for example, 9.3.149.0). The *af* parameter specifies the family of the address. The *src* parameter points to the string being passed in. The *dst* parameter points to a buffer where the function will store the resulting numeric address. The *size* parameter is the size (in bytes) of the buffer pointed to by the *dst* parameter.

## Parameters

<i>af</i>	Specifies the family of the address.
<i>src</i>	Points to the string being passed in.
<i>dst</i>	Points to a buffer where the resulting numeric address is stored.
<i>size</i>	Specifies the size (in bytes) of the buffer pointed to by the <i>dst</i> parameter.

## Return Values

If successful, the number of bits, either inputted classfully or specified with */CIDR*, is returned. If unsuccessful, a -1 (negative one) is returned (check *errno*). *ENOENT* means it was not a valid network specification.

## Related Information

“*inet\_net\_ntop* Subroutine” on page 138, “*inet\_ntop* Subroutine” on page 144, and “*inet\_pton* Subroutine” on page 145.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## inet\_netof Subroutine

### Purpose

Returns the network id of the given Internet address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <sys/socketvar.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
int inet_netof ( InternetAddr)
struct in_addr InternetAddr;
```

### Description

The **inet\_netof** subroutine returns the network number from the specified Internet address number typed as unsigned long value. The **inet\_netof** subroutine masks off the network number and the host number from the Internet address based on the Internet address class.

All applications containing the **inet\_netof** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

**Note:** The application assumes responsibility for verifying that the network number and the host number for the Internet address conforms to a class A or B or C Internet address. The **inet\_netof** subroutine processes any other number as a class C address.

## Parameters

*InternetAddr*

Specifies the Internet address to separate.

## Return Values

Upon successful completion, the **inet\_netof** subroutine returns a network number from the specified long value representing the Internet address. If the application gives a class A Internet address, the **inet\_inoaf** subroutine returns the logical product of the Internet address and 0xFF000000. If the application gives a class B Internet address, the **inet\_inoaf** subroutine returns the logical product of the Internet address and 0xFFFF0000. If the application does not give a class A or B Internet address, the **inet\_inoaf** subroutine returns the logical product of the Internet address and 0xFFFFF000.

## Files

*/etc/hosts*

Contains host names.

*/etc/networks*

Contains network names.

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_makeaddr** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_network Subroutine

### Purpose

Converts an ASCII string containing an Internet network addressee in . (dot) notation to an Internet address number.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
in_addr_t inet_network ( CharString)
register const char *CharString;
```

### Description

The **inet\_network** subroutine converts an ASCII string containing a valid Internet address using . (dot) notation (such as, 120.121.122.123) to an Internet address number formatted as an unsigned integer

value. The **inet\_network** subroutine returns an error value if the application does not provide an ASCII string containing a valid Internet address using . notation.

The input ASCII string must represent a valid Internet address number, as described in "TCP/IP addressing" in *Networks and communication management*. The input string must be terminated with a null terminator (0x00) or a space (0x30). The **inet\_network** subroutine ignores characters that follow the terminating character.

The input string can express an Internet address number in decimal, hexadecimal, or octal format. In hexadecimal format, the string must begin with 0x. The string must begin with 0 to indicate octal format. In decimal format, the string requires no prefix.

Each octet of the input string must be delimited from another by a period. The application can omit values between delimiters. The **inet\_network** subroutine interprets missing values as 0.

The following examples show valid strings and their output values in both decimal and hexadecimal notation:

*Examples of valid strings*

Input String	Output Value (in decimal)	Output Value (in hex)
...1	1	0x00000001
.1..	65536	0x00010000
1	1	0x1
0xFFFFFFFF	255	0x000000FF
1.	256	0x100
1.2.3.4	66048	0x010200
0x01.0x2.03.004	16909060	0x01020304
1.2. 3.4	16777218	0x01000002
9999.1.1.1	251724033	0x0F010101

The following examples show invalid input strings and the reasons they are not valid:

*Examples of invalid strings*

Input String	Reason
1.2.3.4.5	Excessive fields.
1.2.3.4.	Excessive delimiters (and therefore fields).
1,2	Bad delimiter.
1p	String not terminated by null terminator nor space.
{empty string}	No field or delimiter present.

Typically, the value of each octet of an Internet address cannot exceed 246. The **inet\_network** subroutine can accept larger values, but it uses only the eight least significant bits for each field value. For example, if an application passes 0x1234567890.0xabcdef, the **inet\_network** subroutine returns 37103 (0x000090EF).

The application must verify that the network ID and host ID for the Internet address conform to class A, class B, or class C. The **inet\_makeaddr** subroutine processes any nonconforming number as a class C address.

The **inet\_network** subroutine does not check the pointer to the ASCII input string. The application must verify the validity of the address of the string.

All applications containing the **inet\_network** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*CharString*                Represents a string of characters in the Internet address form.

## Return Values

For valid input strings, the **inet\_network** subroutine returns an unsigned integer value that comprises the bit patterns of the input fields concatenated together. The **inet\_network** subroutine places the first pattern in the leftmost (most significant) position and appends subsequent patterns if they exist.

For invalid input strings, the **inet\_network** subroutine returns a value of -1.

## Files

**/etc/hosts**

Contains host names.

**/etc/networks**

Contains network names.

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_makeaddr** subroutine, **inet\_netof** subroutine, **inet\_ntoa** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_ntoa Subroutine

### Purpose

Converts an Internet address into an ASCII string.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>
```

```
char *inet_ntoa ( InternetAddr)
struct in_addr InternetAddr;
```

### Description

The **inet\_ntoa** subroutine takes an Internet address and returns an ASCII string representing the Internet address in dot notation. All Internet addresses are returned in network order, with the first byte being the high-order byte.

Use C language integers when specifying each part of a dot notation.

All applications containing the **inet\_ntoa** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*InternetAddr*                      Contains the Internet address to be converted to ASCII.

## Return Values

Upon successful completion, the **inet\_ntoa** subroutine returns an Internet address.

If the **inet\_ntoa** subroutine is unsuccessful, the subroutine returns a -1.

## Files

*/etc/hosts*    Contains host names.

*/etc/networks*    Contains network names.

## Related Information

The **endhostent** subroutine, **endnetent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_makeaddr** subroutine, **inet\_network** subroutine, **sethostent** subroutine, **setnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## inet\_ntop Subroutine

### Purpose

Converts a binary address into a text string suitable for presentation.

### Library

Library (**libc.a**)

### Syntax

```
const char *inet_ntop (af, src, dst, size)
int af;
const void *src;
char *dst;
size_t size;
```

### Description

This function converts from an address in binary format (as specified by the *src* parameter) to standard text format, and places the result in the *dst* parameter (if *size*, which specifies the space available in the *dst* parameter, is sufficient). The *af* parameter specifies the family of the address. This can be AF\_INET or AF\_INET6.

The *src* parameter points to a buffer holding an IPv4 address if the *af* parameter is AF\_INET, or an IPv6 address if the *af* parameter is AF\_INET6. The *dst* parameter points to a buffer where the function will store the resulting text string. The *size* parameter specifies the size of this buffer (in bytes). The application must specify a non-NULL *dst* parameter. For IPv6 addresses, the buffer must be at least INET6\_ADDRSTRLEN bytes. For IPv4 addresses, the buffer must be at least INET\_ADDRSTRLEN bytes.

In order to allow applications to easily declare buffers of the proper size to store IPv4 and IPv6 addresses in string form, the following two constants are defined in the `<netinet/in.h>` library:

```
#define INET_ADDRSTRLEN 16
#define INET6_ADDRSTRLEN 46
```

## Parameters

<i>af</i>	Specifies the family of the address. This can be AF_INET or AF_INET6.
<i>src</i>	Points to a buffer holding an IPv4 address if the <i>af</i> parameter is set to AF_INET, or an IPv6 address if the <i>af</i> parameter is set to AF_INET6.
<i>dst</i>	Points to a buffer where the resulting text string is stored.
<i>size</i>	Specifies the size (in bytes) of the buffer pointed to by the <i>dst</i> parameter.

## Return Values

If successful, a pointer to the buffer containing the converted address is returned. If unsuccessful, NULL is returned. Upon failure, the **errno** global variable is set to EAFNOSUPPORT if the specified address family (*af*) is unsupported, or to ENOSPC if the *size* parameter indicates the destination buffer is too small.

## Related Information

“inet\_net\_ntop Subroutine” on page 138, “inet\_net\_pton Subroutine” on page 139, and “inet\_pton Subroutine.”

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## inet\_pton Subroutine

### Purpose

Converts an address in its standard text form into its numeric binary form.

### Library

Library (**libc.a**)

### Syntax

```
int inet_pton (af, src, dst)
int af;
const char *src;
void *dst;
```

### Description

This function converts an address in its standard text format into its numeric binary form. The *af* parameter specifies the family of the address.

**Note:** Only the AF\_INET and AF\_INET6 address families are supported.

## Parameters

<i>af</i>	Specifies the family of the address. This can be AF_INET or AF_INET6.
<i>src</i>	Points to the string being passed in.
<i>dst</i>	Points to a buffer where the function stores the numeric address. The address is returned in network byte order.

## Return Values

If successful, one is returned. If unsuccessful, zero is returned if the input is not a valid IPv4 dotted-decimal string or a valid IPv6 address string; or a negative one with the **errno** global variable set to EAFNOSUPPORT if the *af* parameter is unknown. The calling application must ensure that the buffer referred to by the *dst* parameter is large enough to hold the numeric address (4 bytes for AF\_INET or 16 bytes for AF\_INET6).

If the *af* parameter is AF\_INET, the function accepts a string in the standard IPv4 dotted-decimal form.

*ddd.ddd.ddd.ddd*

Where *ddd* is a one to three digit decimal number between 0 and 255.

**Note:** Many implementations of the existing **inet\_addr** and **inet\_aton** functions accept nonstandard input such as octal numbers, hexadecimal numbers, and fewer than four numbers. **inet\_pton** does not accept these formats.

If the *af* parameter is AF\_INET6, then the function accepts a string in one of the standard IPv6 text forms defined in the addressing architecture specification.

## Related Information

“inet\_net\_ntop Subroutine” on page 138, “inet\_net\_pton Subroutine” on page 139, and “inet\_ntop Subroutine” on page 144.

**Subroutines Overview** in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

---

## inetgr, getnetgrent, setnetgrent, or endnetgrent Subroutine

### Purpose

Handles the group network entries.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
```

```
inetgr (NetGroup, Machine, User, Domain)
char * NetGroup, * Machine, * User, * Domain;
```

```
getnetgrent (MachinePointer, UserPointer, DomainPointer)
char ** MachinePointer, ** UserPointer, ** DomainPointer;
void setnetgrent (NetGroup)
char *NetGroup
void endnetgrent ()
```

## Description

The **innetgr** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **innetgr** subroutine returns *1* or *0*, depending on if **netgroup** contains the *machine*, *user*, *domain* triple as a member. Any of these three strings; *machine*, *user*, or *domain*, can be NULL, in which case it signifies a wild card.

The **getnetgrent** subroutine returns the next member of a network group. After the call, *machinepointer* will contain a pointer to a string containing the name of the machine part of the network group member, and similarly for *userpointer* and *domainpointer*. If any of *machinepointer*, *userpointer*, or *domainpointer* is returned as a NULL pointer, it signifies a wild card. The **getnetgrent** subroutine uses malloc to allocate space for the name. This space is released when the **endnetgrent** subroutine is called. **getnetgrent** returns *1* if it succeeded in obtaining another member of the network group or *0* when it has reached the end of the group.

The **setnetgrent** subroutine establishes the network group from which the **getnetgrent** subroutine will obtain members, and also restarts calls to the **getnetgrent** subroutine from the beginning of the list. If the previous **setnetgrent()** call was to a different network group, an **endnetgrent()** call is implied. **endnetgrent()** frees the space allocated during the **getnetgrent()** calls.

## Parameters

<i>Domain</i>	Specifies the domain.
<i>DomainPointer</i>	Points to the string containing <i>Domain</i> part of the network group.
<i>Machine</i>	Specifies the machine.
<i>MachinePointer</i>	Points to the string containing <i>Machine</i> part of the network group.
<i>NetGroup</i>	Points to a network group.
<i>User</i>	Specifies a user.
<i>UserPointer</i>	Points to the string containing <i>User</i> part of the network group.

## Return Values

- 1** Indicates that the subroutine was successful in obtaining a member.
- 0** Indicates that the subroutine was not successful in obtaining a member.

## Files

<i>/etc/netgroup</i>	Contains network groups recognized by the system.
<i>/usr/include/netdb.h</i>	Contains the network database structures.

## Related Information

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## ioctl Socket Control Operations

### Purpose

Performs network related control operations.

## Syntax

```
#include <sys/ioctl.h>

int ioctl (fd, cmd, .../* arg */)
int fd;
int cmd;
int ... /* arg */
```

## Description

The socket `ioctl` commands perform a variety of network-related control. The *fd* argument is a socket descriptor. For non-socket descriptors, the functions performed by this call are unspecified.

The *cmd* argument and an optional third argument (with varying type) are passed to and interpreted by the socket `ioctl` function to perform an appropriate control operation specified by the user.

The socket `ioctl` control operations can be in the following control operations categories:

- Socket
- Routing table
- ARP table
- Global network parameters
- Interface

## Parameters

<i>fd</i>	Open file descriptor that refers to a socket created using <code>socket</code> or <code>accept</code> calls.
<i>cmd</i>	Selects the control function to be performed.
<i>.../* arg */</i>	Represents additional information that is needed to perform the requested function. The type of <i>arg</i> depends on the particular control request, but it is either an integer or a pointer to a socket-specific data structure.

## Socket Control Operations

The following `ioctl` commands operate on sockets:

ioctl command	Description
SIOCATMARK	Determines whether the read pointer is currently pointing to the logical mark in the data stream. The logical mark indicates the point at which the out-of-band data is sent.  <code>ioctl(fd, SIOCATMARK, &amp;atmark);</code> <code>int atmark;</code>  If <i>atmark</i> is set to 1 on return, the read pointer points to the mark and the next read returns data after the mark. If <i>atmark</i> is set to 0 on return, (assuming out-of-band data is present on the data stream), the next read returns data sent prior to the out-of-band mark. <b>Note:</b> The out-of-band data is a logically independent data channel that is delivered to the user independently of normal data; in addition, a signal is also sent because of the immediate attention required. Ctrl-C characters are an example.
SIOCSPGRP SIOCGPGRP	SIOCSPGRP sets the process group information for a socket. SIOCGPGRP gets the process group ID associated with a socket.  <code>ioctl(fd, cmd, (int)&amp;pgrp);</code> <code>int pgrp;</code>  <i>cmd</i> Set to SIOCSPGRP or SIOCGPGRP.  <i>pgrp</i> Specifies the process group ID for the socket.

## Routing Table Control Operations

The following ioctl commands operate on the kernel routing table:

ioctl command	Description
SIOCADDRT SIOCDELRT	<p>SIOCADDRT adds a route entry in the routing table. SIOCDELRT deletes a route entry from the routing table.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;route); struct ortentry route;</pre> <p><i>cmd</i> Set to SIOCADDRT or SIOCDELRT.</p> <p>The route entry information is passed in the <b>ortentry</b> structure.</p>
SIOUPDRROUTE	<p>Updates the routing table using the information passed in the <b>ifreq</b> structure.</p> <pre>ioctl(fd, SIOUPDRROUTE, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p><b>Note:</b> SIOUPDRROUTE is available beginning with AIX 5100-002 (maintenance level 2) and later.</p>

## ARP Table Control Operations

The following ioctl commands operate on the kernel ARP table. The **net/if\_arp.h** header file must be included.

ioctl command	Description
SIOCSARP SIOCDAARP SIOCGARP	<p>SIOCSARP adds or modifies an ARP entry in the ARP table. SIOCDAARP deletes an ARP entry from the ARP table. SIOCGARP gets an ARP entry from the ARP table.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ar); struct arpreq ar;</pre> <p><i>cmd</i> Set to SIOCSARP, SIOCDAARP, or SIOCGARP.</p> <p>The ARP entry information is passed in the <b>arpreq</b> structure.</p>

## Global Network Parameters Control Operations

The following ioctl commands operate as global network parameters:

ioctl command	Description
SIOCSNETOPT SIOCGNETOPT SIOCNETOPT SIOCGNETOPT1	<p>SIOCSNETOPT sets the value of a network option. SIOCGNETOPT gets the value of a network option. SIOCNETOPT sets the default values of a network option.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;oreq); struct optreq oreq;</pre> <p><i>cmd</i> Set to SIOCSNETOPT, SIOCGNETOPT, or SIOCNETOPT.</p> <p>The network option value is stored in the <b>optreq</b> structure.</p> <p>SIOCGNETOPT1 gets the current value, default value, and the range of a network option.</p> <pre>ioctl(fd, SIOCGNETOPT1, (caddr_t)&amp;oreq); struct optreq1 oreq1;</pre> <p>The network option information is stored in the <b>optreq1</b> structure upon return. The <b>optreq</b> and <b>optreq1</b> structures are defined in <b>net/netopt.h</b>.</p> <p><b>Note:</b> SIOCGNETOPT1 is available beginning with AIX 5.2 and later.</p>

ioctl command	Description
SIOCGNMTUS SIOCGETMTUS SIOCADDMTU SIOCDELMTU	<p>SIOCGNMTUS gets the number of MTUs maintained in the list of common MTUs. SIOCADDMTU adds an MTU in the list of common MTUs. SIOCDELMTU deletes an MTU from the list of common MTUs.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;nmtus); int nmtus;</pre> <p><i>cmd</i> Set to SIOCGNMTUS, SIOCADDMTU, or SIOCDELMTU.</p> <p>SIOCGETMTUS gets the MTUs maintained in the list of common MTUs.</p> <pre>ioctl(fd, SIOCGETMTUS, (caddr_t)&amp;gm); struct get_mtus gm;</pre> <p>The <b>get_mtus</b> structure is defined in <b>netinet/in.h</b>.</p>

## Interface Control Operations

The following ioctl commands operate on interfaces. The **net/if.h** header file must be included.

ioctl command	Description
SIOCSIFADDR SIOCAIFADDR SIOCDEFADDR SIOCGIFADDR	<p>SIOCSIFADDR sets an interface address. SIOCAIFADDR adds or changes an interface address. SIOCDEFADDR deletes an interface address. The interface address is specified in the <i>ifr.ifr_addr</i> field. SIOCGIFADDR gets an interface address. The address is returned in the <i>ifr.ifr_addr</i> field.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr, sizeof(struct ifreq)); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSIFADDR, SIOCAIFADDR, SIOCDEFADDR, or SIOCGIFADDR.</p>
SIOCGIFADDRS	<p>Gets the list of addresses associated with an interface.</p> <pre>ioctl(fd, SIOCGIFADDRS, (caddr_t)ifaddrsp); struct ifreqaddrs *ifaddrsp;</pre> <p>The interface name is passed in the <i>ifaddrsp-&gt;ifr_name</i> field. The addresses associated with the interface are stored in <i>ifaddrsp-&gt;ifrasu[]</i> array on return.</p> <p><b>Note:</b> The <b>ifreqaddrs</b> structure contains space for storing only one <b>sockaddr_in/sockaddr_in6</b> structure (array of one <b>sockaddr_in/sockaddr_in6</b> element). In order to get <i>n</i> addresses associated with an interface, the caller of the ioctl command must allocate space for <i>{sizeof (struct ifreqaddrs) + (n * sizeof (struct sockaddr_in))}</i> bytes.</p> <p><b>Note:</b> SIOCGIFADDRS is available beginning with AIX 5.3 and later.</p>
SIOCSIFDSTADDR SIOCGIFDSTADDR	<p>SIOCSIFDSTADDR sets the point-to-point address for an interface specified in the <i>ifr.ifr_dstaddr</i> field. SIOCGIFDSTADDR gets the point-to-point address associated with an interface. The address is stored in the <i>ifr.ifr_dstaddr</i> field on return.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr, sizeof(struct ifreq)); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSIFDSTADDR or SIOCGIFDSTADDR.</p>
SIOCSIFNETMASK SIOCGIFNETMASK	<p>SIOCSIFNETMASK sets the interface netmask specified in the <i>ifr.ifr_addr</i> field. SIOCGIFNETMASK gets the interface netmask.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr, sizeof(struct ifreq)); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSIFNETMASK or SIOCGIFNETMASK.</p>

ioctl command	Description
SIOCSIFBRDADDR SIOCGIFBRDADDR	<p>SIOCSIFBRDADDR sets the interface broadcast address specified in the <i>ifr.ifr_broadaddr</i> field. SIOCGIFBRDADDR gets the interface broadcast address. The broadcast address is placed in the <i>ifr.ifr_broadaddr</i> field.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr, sizeof(struct ifreq)); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSIFBRDADDR or SIOCGIFBRDADDR.</p>
SIOCGSIZIFCONF	<p>Gets the size of memory required to get configuration information for all interfaces returned by SIOCGIFCONF.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifconfsize); int ifconfsize;</pre>
SIOCGIFCONF	<p>Returns configuration information for all the interfaces configured on the system.</p> <pre>ioctl(fd, SIOCGIFCONF, (caddr_t)&amp;ifc); struct ifconf ifc;</pre> <p>The configuration information is returned in a list of <b>ifreq</b> structures pointed to by the <i>ifc.ifc_req</i> field, with one <b>ifreq</b> structure per interface.</p> <p><b>Note:</b> The caller of the ioctl command must allocate sufficient space to store the configuration information, returned as a list of <b>ifreq</b> structures for all of the interfaces configured on the system. For example, if <i>n</i> interfaces are configured on the system, <i>ifc.ifc_req</i> must point to <i>{n * sizeof (struct ifreq)}</i> bytes of space allocated.</p> <p><b>Note:</b> Alternatively, the SIOCGSIZIFCONF ioctl command can be used for this purpose.</p>
SIOCSIFFLAGS SIOCGIFFLAGS	<p>SIOCSIFFLAGS sets the interface flags. SIOCGIFFLAGS gets the interface flags.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p>Refer to <code>/usr/include/net/if.h</code> for the interface flags, denoted by <b>IFF_</b>xxx.</p> <p><b>Note:</b> The <b>IFF_BROADCAST</b>, <b>IFF_POINTTOPPOINT</b>, <b>IFF_SIMPLEX</b>, <b>IFF_RUNNING</b>, <b>IFF_OACTIVE</b>, and <b>IFF_MULTICAST</b> flags cannot be changed using ioctl.</p>
SIOCSIFMETRIC SIOCGIFMETRIC	<p>SIOCSIFMETRIC sets the interface metric specified in the <i>ifr.ifr_metric</i> field. SIOCGIFMETRIC gets the interface metric. The interface metric is placed in the <i>ifr.ifr_metric</i> field on return.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSIFMETRIC or SIOCGIFMETRIC.</p>
SIOCSIFSUBCHAN SIOCGIFSUBCHAN	<p>SIOCSIFSUBCHAN sets the subchannel address specified in the <i>ifr.ifr_flags</i> field. SIOCGIFSUBCHAN gets the subchannel address in the <i>ifr.ifr_flags</i> field.</p> <pre>ioctl(fd, SIOCSIFSUBCHAN, (caddr_t)&amp;ifr); struct ifreq ifr;</pre>
SIOCSIFOPTIONS SIOCGIFOPTIONS	<p>SIOCSIFOPTIONS sets the interface options. SIOCGIFOPTIONS gets the interface options.</p> <pre>ioctl(fd, SIOCSIFOPTIONS, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p>The interface options are stored in the <i>ifr_flags</i> field of the <b>ifreq</b> structure. Refer to <code>/usr/include/net/if.h</code> file for the list of interface options denoted by <b>IFO_</b>xxx.</p>

ioctl command	Description
SIOCADDMULTI SIOCDELMULTI	<p>SIOCADDMULTI adds an address to the list of multicast addresses for an interface. SIOCDELMULTI deletes a multicast address from the list of multicast addresses for an interface.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCADDMULTI or SIOCDELMULTI.</p> <p>The multicast address information is specified in the <b>ifr_addr</b> structure.</p>
SIOCGETVIFCNT	<p>Gets the packet count information for a virtual interface. The information is specified in the <b>sioc_vif_req</b> structure.</p> <pre>ioctl (fd, SIOCGETVIFCNT, (caddr_t)&amp;v_req); struct sioc_vif_req v_req;</pre>
SIOCGETSGCNT	<p>Gets the packet count information for the source group specified. The information is stored in the <b>sioc_sg_req</b> structure on return.</p> <pre>ioctl(fd, SIOCGETSGCNT, (caddr_t)&amp;v_req); struct sioc_sg_req v_req;</pre>
SIOCSIFMTU SIOCGIFMTU	<p>SIOCSIFMTU sets the interface maximum transmission unit (MTU). SIOCGIFMTU gets the interface MTU.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p>The MTU value is stored in <i>ifr.ifr_mtu</i> field.</p> <p><b>Note:</b> The range of valid values for MTU varies for an interface depending on the interface type.</p>
SIOCIFATTACH SIOCIFDETACH	<p>SIOCIFATTACH attaches an interface. This initializes and adds an interface in the network interface list. SIOCIFDETACH detaches an interface broadcast address. This removes the interface from the network interface list. The interface name is specified in the <i>ifr.ifr_name</i> field.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre>
SIOCSIFGIDLIST SIOCGIFGIDLIST	<p>SIOCSIFGIDLIST adds or deletes the list of group IDs specified in the <i>ifrg.ifrg_gidlist</i> field to the <i>gidlist</i> interface. The interface name is specified in the <i>ifrg.ifrg_name</i> field. An operation code, ADD_GRP/DEL_GRP, specified in the <i>ifrg.ifrg_gidlist</i> field indicates whether the specified list of group IDs must be added to or deleted from the <i>gidlist</i> interface. SIOCGIFGIDLIST gets the list of group IDs associated with an interface. The group IDs are placed in the <i>ifrg.ifrg_gidlist</i> field on return.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifrg); struct ifgidreq ifrg;</pre>
SIOCIF_ATM_UBR SIOCIF_ATM_SNMPARP SIOCIF_ATM_DUMPARP SIOCIF_ATM_IDLE SIOCIF_ATM_SVC SIOCIF_ATM_DARP SIOCIF_ATM_GARP SIOCIF_ATM_SARP	<p>SIOCIF_ATM_UBR sets the UBR rate for an ATM interface.</p> <p>SIOCIF_ATM_SNMPARP gets the SNMP ATM ARP entries.</p> <p>SIOCIF_ATM_DUMPARP gets the specified number of ATM ARP entries.</p> <p>SIOCIF_ATM_DARP deletes an ATM ARP entry from the ARP table.</p> <p>SIOCIF_ATM_GARP gets an ATM ARP entry to the ARP table.</p> <p>SIOCIF_ATM_SARP adds an ATM ARP entry. The ARP information is specified in the <b>atm_arpreq</b> structure.</p> <p>SIOCIF_ATM_SVC specifies whether this interface supports Permanent Virtual Circuit (PVC) and Switched Virtual Circuit (SVC) types of virtual connections. It also specifies whether this interface will be an ARP client or an ARP server for this Logical IP Subnetwork (LIS) based on the flag set in the <b>ifatm_svc_arg</b> structure.</p> <p>SIOCIF_ATM_IDLE specifies the idle time limit on the interface.</p>

ioctl command	Description
SIOCSISNO SIOCGISNO	<p>SIOCSISNO sets interface specific network options for an interface. SIOCGISNO gets interface specific network options associated with an interface.</p> <pre>ioctl(fd, cmd, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p><i>cmd</i> Set to SIOCSISNO or SIOCGISNO.</p> <p>The interface specific network options are stored in <b>ifr.ifr_isno</b> structure. Refer to <b>/usr/include/net/if.h</b> file for the list of interface specific network options denoted by <b>ISNO_XXX</b>.</p>
SIOCGIFBAUDRATE	<p>Gets the value of the interface baud rate in the <i>ifr_baudrate</i> field.</p> <pre>ioctl(fd, SIOCGIFBAUDRATE, (caddr_t)&amp;ifr); struct ifreq ifr;</pre> <p>The baud rate is stored in the <i>ifr.ifr_baudrate</i> field.</p>
SIOCADDIFVIPA SIOCDELIFVIPA SIOCLISTIFVIPA	<p>SIOCADDIFVIPA associates the specified list of interfaces pointed to by <i>ifrv.ifrv_ifname</i> with the virtual interface specified by <i>ifrv.ifrv_name</i>. This causes the source address for all outgoing packets on these interfaces to be set to the virtual interface address. SIOCDELIFVIPA removes the list of specified interfaces pointed to by <i>ifrv.ifrv_ifname</i>, that are associated with the virtual interface specified by <i>ifrv.ifrv_name</i>, using SIOCADDIFVIPA. SIOCLISTIFVIPA lists all the interfaces associated with the virtual interface specified by <i>ifrv.ifrv_name</i>.</p> <pre>ioctl(fd, SIOCADDIFVIPA, (caddr_t)&amp;ifrv); struct ifvireq ifrv;</pre> <p>The virtual interface information is stored in the <b>ifvireq</b> structure.</p> <p><b>Note:</b> These flags operate on a virtual interface only. Also, these flags are available beginning with AIX 5.2 and later.</p>

## Return Values

Upon successful completion, ioctl returns 0. Otherwise, it returns -1 and sets **errno** to indicate the error.

## Error Codes

The ioctl commands fail under the following general conditions:

<b>EBADF</b>	The file descriptor <i>fd</i> is not a valid open socket file descriptor.
<b>EINTR</b>	A signal was caught during ioctl operation.
<b>EINVAL</b>	An invalid command or argument was specified.

If the underlying operation specified by the ioctl command *cmd* failed, ioctl fails with one of the following error codes:

<b>EACCES</b>	Permission denied for the specified operation.
<b>EADDRNOTAVAIL</b>	Specified address not available for interface.
<b>EAFNOSUPPORT</b>	Operation not supported on sockets.
<b>EBUSY</b>	Resource is busy.
<b>EEXIST</b>	An entry or file already exists.
<b>EFAULT</b>	Argument references an inaccessible memory area.
<b>EIO</b>	I/O error.
<b>ENETUNREACH</b>	Gateway unreachable.
<b>ENOBUFS</b>	Routing table overflow.
<b>ENOCCONNECT</b>	No connection.

<b>ENOMEM</b>	Not enough memory available.
<b>ENOTCONN</b>	The operation is only defined on a connected socket, but the socket was not connected.
<b>ENXIO</b>	Device does not exist.
<b>ESRCH</b>	No such process.

## Related Information

Socket Overview.

The `ioctl` subroutine.

“`ioctl` Streams Device Driver Operations” on page 281.

---

## isinet\_addr Subroutine

### Purpose

Determines if the given ASCII string contains an Internet address using dot notation.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
u_long isinet_addr (name)
char * name;
```

### Description

The **isinet\_addr** subroutine determines if the given ASCII string contains an Internet address using dot notation (for example, “120.121.122.123”). The **isaddr\_inet** subroutine considers Internet address strings as a valid string, and considers any other string type as an invalid strings.

The **isinet\_addr** subroutine expects the ASCII string to conform to the following format:

```
string ::= field | field delimited_field^1-3
delimited_field ::= delimiter field
delimiter ::= .
field ::= 0 x | 0 x | 0 x hexadecimal* | 0 x hexadecimal* | decimal* | 0 octal*
hexadecimal ::= decimal | a | b | c | d | e | f | A | B | C | D | E | F
decimal ::= octal | 8 | 9
octal ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7
```

Value	Description
$A^n$	Indicates $n$ repetitions of pattern A.
$A^n-m$	Indicates $n$ to $m$ repetitions of pattern A.
$A^*$	Indicates zero or more repetitions of pattern A, up to environmental limits.

The BNF description explicitly states the space character (‘ ’), if used.

Value	Description
{ <i>text</i> }	Indicates <i>text</i> , not a BNF symbol.

The **isinet\_addr** subroutine allows the application to terminate the string with a null terminator (0x00) or a space (0x30). It ignores characters trailing the space character and considers the string invalid if the application does not terminate the string with a null terminator (0x00) or space (0x30).

The following describes the restrictions on the field values:

Address Format	Field Restrictions (values in decimal base)
a	a < 4294967296.
a.b	a < 256; b < 16777216.
a.b.c	a < 256; b < 256; c < 16777216.
a.b.c.d	a < 256; b < 2 <sup>8</sup> ; c < 256; d < 256.

The **isinet\_addr** subroutine applications can enter field values exceeding the field value restrictions specified previously; **isinet\_addr** accepts the least significant bits up to an integer in length. The **isinet\_addr** subroutine still checks to see if the truncated value exceeds the maximum field value. For example, if an application gives the string 0.0;0;0xFF00000001 then **isinet\_addr** interprets the string as 0.0.0.0x00000001 and considers the string as valid.

**isinet\_addr** applications cannot omit field values between delimiters and considers a string with successive periods as invalid.

Examples of valid strings:

Input String	Comment
1	<b>isinet_addr</b> uses <b>a</b> format.
1.2	<b>isinet_addr</b> uses <b>a.b</b> format.
1.2.3.4	<b>isinet_addr</b> uses <b>a.b.c.d</b> format.
0x01.0x2.03.004	<b>isinet_addr</b> uses <b>a.b.c.d</b> format.
1.2 3.4	<b>isinet_addr</b> uses <b>a.b</b> format; and ignores "3.4".

Examples of invalid strings:

Input String	Reason
...	No explicit field values specified.
1.2.3.4.5	Excessive fields.
1.2.3.4.	Excessive delimiters and fields.
1,2	Bad delimiter.
1p	String not terminated by null terminator nor space.
{empty string}	No field or delimiter present.
9999.1.1.1	Value for field a exceeds limit.

#### Notes:

1. The **isinet\_addr** subroutine does not check the pointer to the ASCII string; the user takes responsibility for ensuring validity of the address of the ASCII string.
2. The application assumes responsibility for verifying that the network number and host number for the Internet address conforms to a class A or B or C Internet address; any other string is processed as a class C address.

All applications using **isinet\_addr** must compile with the **\_BSD** macro defined. Also, all socket applications must include the BSD library **libbsd** when applicable.

## Parameters

*name*    Address of ASCII string buffer.

## Return Values

The **isinet\_addr** subroutine returns 1 for valid input strings and 0 for invalid input strings. **isinet\_addr** returns the value as an unsigned long type.

## Files

**#include <ctype.h>**

**#include <sys/types.h>**

## Related Information

Internet address conversion subroutines: **inet\_addr** subroutine, **inet\_lnaof** subroutine, **inet\_makeaddr** subroutine, **inet\_netof** subroutine, **inet\_network** subroutine, **inet\_ntoa** subroutine.

Host information retrieval subroutines: **endhostent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **sethostent** subroutine.

Network information retrieval subroutines: **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine, **setnetent** subroutine.

---

## kvalid\_user Subroutine

### Purpose

This routine maps the DCE principal to the local user account and determines if the DCE principal is allowed access to the account.

### Library

Valid User Library (**libvaliduser.a**)

### Syntax

### Description

This routine is called when Kerberos 5 authentication is configured to determine if the incoming Kerberos 5 ticket should allow access to the local account.

This routine determines whether the DCE principal, specified by the *princ\_name* parameter, is allowed access to the user's account identified by the *local\_user* parameter. The routine accesses the **\$HOME/.k5login** file for the user's account. It looks for the string pointed to by *princ\_name* in that file.

Access is granted if one of two things is true.

1. The **\$HOME/.k5login** file exists and the *princ\_name* is in it.
2. The **\$HOME/.k5login** file does NOT exist and the DCE principal name is the same as the local user's name.

### Parameters

<i>princ_name</i>	This parameter is a single-string representation of the Kerberos 5 principal. The Kerberos 5 libraries have two services, <i>krb5_unparse_name</i> and <i>krb5_parse_name</i> , which convert a <i>krb5_principal</i> structure to and from a single-string format. This routine expects the <i>princ_name</i> parameter to be a single-string form of the <i>krb5_principal</i> structure.
-------------------	---

*local\_user* This parameter is the character string holding the name of the local account.

## Return Values

If the user is allowed access to the account, the **kvalid\_user** routine returns TRUE.

If the user is NOT allowed access to the account or there was an error, the **kvalid\_user** routine returns FALSE.

## Related Information

The **ftp** command, **rcp** command, **rlogin** command, **rsh** command, **telnet**, **tn**, or **tn3270** command.

Using a **.k5login** file.

Communications and networks in *Networks and communication management*.

Authentication and the secure rcmds in *Networks and communication management*.

---

## listen Subroutine

### Purpose

Listens for socket connections and limits the backlog of incoming connections.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int listen ( Socket, Backlog)
int Socket, Backlog;
```

### Description

The **listen** subroutine performs the following activities:

1. Identifies the socket that receives the connections.
2. Marks the socket as accepting connections.
3. Limits the number of outstanding connection requests in the system queue.

The outstanding connection request queue length limit is specified by the parameter *backlog* per listen call. A *no* parameter - *somaxconn* - defines the maximum queue length limit allowed on the system, so the effective queue length limit will be either *backlog* or *somaxconn*, whichever is smaller.

All applications containing the **listen** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the unique name for the socket.
<i>Backlog</i>	Specifies the maximum number of outstanding connection requests.

## Return Values

Upon successful completion, the **listen** subroutine returns a value 0.

If the **listen** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNREFUSED</b>	The host refused service, usually due to a server process missing at the requested name or the request exceeding the backlog amount.
<b>EINVAL</b>	The socket is already connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The referenced socket is not a type that supports the <b>listen</b> subroutine.

## Examples

The following program fragment illustrates the use of the **listen** subroutine with 5 as the maximum number of outstanding connections which may be queued awaiting acceptance by the server process.

```
listen(s,5)
```

## Related Information

The **accept** subroutine, **connect** subroutine, **socket** subroutine.

Accepting Internet Stream Connections Example Program, Sockets Overview, Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

---

## ntohl Subroutine

### Purpose

Converts an unsigned long integer from Internet network standard byte order to host byte order.

### Library

ISODE Library (**libisode.a**)

### Syntax

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

```
uint32_t ntohl ( NetLong)
uint32_t NetLong;
```

### Description

The **ntohl** subroutine converts an unsigned long (32-bit) integer from Internet network standard byte order to host byte order.

Receiving hosts require addresses and ports in host byte order. Use the **ntohl** subroutine to convert Internet addresses and ports to the host integer representation.

The **ntohl** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order.

The **ntohl** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **ntohl** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*NetLong*      Requires a 32-bit integer in network byte order.

## Return Values

The **ntohl** subroutine returns a 32-bit integer in host byte order.

## Related Information

The **endhostent** subroutine, **endservent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **getservent** subroutine, **htonl** subroutine, **htons** subroutine, **ntohs** subroutine, **htonll** subroutine, **ntohll** subroutine, **sethostent** subroutine, **setservent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## ntohll Subroutine

### Purpose

Converts an unsigned long integer from Internet network standard byte order to host byte order.

### Library

ISODE Library (**libisode.a**)

### Syntax

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

```
uint64_t ntohll ( NetLong)
uint64_t NetLong;
```

### Description

The **ntohll** subroutine converts an unsigned long (64-bit) integer from Internet network standard byte order to host byte order.

Receiving hosts require addresses and ports in host byte order. Use the **ntohll** subroutine to convert Internet addresses and ports to the host integer representation.

The **ntohll** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is the same as the network byte order.

The **ntohll** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not the same as the network byte order.

All applications containing the **ntohl** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*NetLong* Requires a 64-bit integer in network byte order.

## Return Values

The **ntohl** subroutine returns a 64-bit integer in host byte order.

## Related Information

The **endhostent** subroutine, **endservent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **getservent** subroutine, **htonl** subroutine, **htons** subroutine, **ntohs** subroutine, **htonll** subroutine, **ntohl** subroutine, **sethostent** subroutine, **setservent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## ntohs Subroutine

### Purpose

Converts an unsigned short integer from Internet network byte order to host byte order.

### Library

ISODE Library (**libisode.a**)

### Syntax

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

```
uint16_t ntohs ( NetShort)
uint16_t NetShort;
```

### Description

The **ntohs** subroutine converts an unsigned short (16-bit) integer from Internet network byte order to the host byte order.

Receiving hosts require Internet addresses and ports in host byte order. Use the **ntohs** subroutine to convert Internet addresses and ports to the host integer representation.

The **ntohs** subroutine is defined in the **net/nh.h** file as a null macro if the host byte order is same as the network byte order.

The **ntohs** subroutine is declared in the **net/nh.h** file as a function if the host byte order is not same as the network byte order.

All applications containing the **ntohs** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*NetShort* Requires a 16-bit integer in network standard byte order.

## Return Values

The **ntohs** subroutine returns a 16-bit integer in host byte order.

## Related Information

The **endhostent** subroutine, **endservent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **getservent** subroutine, **htonl** subroutine, **htons** subroutine, **ntohl** subroutine, **htonll** subroutine, **ntohl** subroutine, **sethostent** subroutine, **setservent** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## PostQueuedCompletionStatus Subroutine

### Purpose

Post a completion packet to a specified I/O completion port.

### Syntax

```
#include <iocp.h>
boolean_t PostQueuedCompletionStatus (CompletionPort, TransferCount, CompletionKey, Overlapped, )
HANDLE CompletionPort;
DWORD TransferCount, CompletionKey;
LPOVERLAPPED Overlapped;
```

### Description

The **PostQueuedCompletionStatus** subroutine attempts to post a completion packet to *CompletionPort* with the values of the completion packet populated by the *TransferCount*, *CompletionKey*, and *Overlapped* parameters.

The **PostQueuedCompletionStatus** subroutine returns a boolean indicating whether or not a completion packet has been posted.

The **PostQueuedCompletionStatus** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

**Note:** This subroutine only works to a socket file descriptor. It does not work with files or other file descriptors.

## Parameters

<i>CompletionPort</i>	Specifies the completion port that this subroutine will attempt to access.
<i>TransferCount</i>	Specifies the number of bytes transferred.
<i>CompletionKey</i>	Specifies the completion key.
<i>Overlapped</i>	Specifies the overlapped structure.

## Return Values

Upon successful completion, the **PostQueuedCompletionStatus** subroutine returns a boolean indicating its success.

If the **PostQueuedCompletionStatus** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of 0 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable, see the link in the Related Information section of this document.

## Error Codes

The subroutine is unsuccessful if either of the following errors occur:

<b>EBADF</b>	The <i>CompletionPort</i> parameter was NULL.
<b>EINVAL</b>	The <i>CompletionPort</i> parameter was invalid.

## Examples

The following program fragment illustrates the use of the **PostQueuedCompletionStatus** subroutine to post a completion packet.

```
c = GetQueuedCompletionStatus (34, 128, 25, struct overlapped);
```

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “ReadFile Subroutine” on page 165, “WriteFile Subroutine” on page 258, “GetQueuedCompletionStatus Subroutine” on page 101, “GetMultipleCompletionStatus Subroutine” on page 83, and “CreateIoCompletionPort Subroutine” on page 37.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*

---

## rcmd Subroutine

### Purpose

Allows execution of commands on a remote host.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int rcmd (Host,  
Port, LocalUser, RemoteUser, Command, ErrFileDesc)  
char ** Host;  
u_short Port;  
char * LocalUser;  
char * RemoteUser;  
char * Command;  
int * ErrFileDesc;
```

### Description

The **rcmd** subroutine allows execution of certain commands on a remote host that supports **rshd**, **rlogin**, and **rpc** among others.

Only processes with an effective user ID of root user can use the **rcmd** subroutine. An authentication scheme based on remote port numbers is used to verify permissions. Ports in the range between 0 and 1023 can only be used by a root user. The application must pass in *Port*, which must be in the range 512 to 1023.

The **rcmd** subroutine looks up a host by way of the name server or if the local name server isn't running, in the **/etc/hosts** file.

If the connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process and given to the remote command as standard input (**stdin**) and standard output (**stdout**).

Always specify the *Host* parameter. If the local domain and remote domain are the same, specifying the domain parts is optional.

All applications containing the **rcmd** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Host</i>	Specifies the name of a remote host that is listed in the <b>/etc/hosts</b> file. If the specified name of the host is not found in this file, the <b>rcmd</b> subroutine is unsuccessful.
<i>Port</i>	Specifies the well-known port to use for the connection. The <b>/etc/services</b> file contains the DARPA Internet services, their ports, and socket types.
<i>LocalUser</i> and <i>RemoteUser</i>	Points to user names that are valid at the local and remote host, respectively. Any valid user name can be given.
<i>Command</i>	Specifies the name of the command to be started at the remote host.
<i>ErrFileDesc</i>	Specifies an integer controlling the set up of communication channels. Integer options are as follows:  <b>Non-zero</b> Indicates an auxiliary channel to a control process is set up, and the <i>ErrFileDesc</i> parameter points to the file descriptor for the channel. The control process provides diagnostic output from the remote command on this channel and also accepts bytes as signal numbers to be forwarded to the process group of the command.  <b>0</b> Indicates the standard error ( <b>stderr</b> ) of the remote command is the same as standard output ( <b>stdout</b> ). No provision is made for sending arbitrary signals to the remote process. However, it is possible to send out-of-band data to the remote command.

## Return Values

Upon successful completion, the **rcmd** subroutine returns a valid socket descriptor.

Upon unsuccessful completion, the **rcmd** subroutine returns a value of -1. The subroutine returns a -1, if the effective user ID of the calling process is not root user or if the subroutine is unsuccessful to resolve the host.

## Files

<b>/etc/services</b>	Contains the service names, ports, and socket type.
----------------------	---

<b>/etc/hosts</b>	Contains host names and their addresses for hosts in a network.
<b>/etc/resolv.conf</b>	Contains the name server and domain name.

## Related Information

The **rlogind** command, **rshd** command.

The **named** daemon.

The **gethostname** subroutine, **rresvport** subroutine, **ruserok** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## rcmd\_af Subroutine

### Purpose

Allows execution of commands on a remote host.

### Syntax

```
int rcmd_af(char **ahost, unsigned short rport,
            const char *locuser, const char *remuser,
            const char *cmd, int *fd2p, int af)
```

### Description

The **rcmd\_af** subroutine allows execution of certain commands on a remote host that supports **rshd**, **rlogin**, and **rpc** among others. It behaves the same as the existing **rcmd()** function, but instead of creating only an AF\_INET TCP socket, it can also create an AF\_INET6 TCP socket. The existing **rcmd()** function cannot transparently use AF\_INET6 sockets because an application would not be prepared to handle AF\_INET6 addresses returned by subroutines such as **getpeername()** on the file descriptor created by **rcmd()**.

Only processes with an effective user ID of root user can use the **rcmd\_af** subroutine. An authentication scheme based on remote port numbers is used to verify permissions. Ports in the range between 0 and 1023 can only be used by a root user.

The **rcmd\_af** subroutine looks up a host by way of the name server or if the local name server is not running, in the **/etc/hosts** file.

If the connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process and given to the remote command as standard input (**stdin**) and standard output (**stdout**).

Always specify the *ahost* parameter. If the local domain and remote domain are the same, specifying the domain parts is optional.

### Parameters

<i>ahost</i>	Specifies the name of a remote host that is listed in the <b>/etc/hosts</b> file. If the specified name of the host is not found in this file, the <b>rcmd_af</b> subroutine is unsuccessful.
<i>rport</i>	Specifies the well-known port to use for the connection. The <b>/etc/services</b> file contains the DARPA Internet services, their ports, and socket types.
<i>locuser</i>	Points to user names that are valid at the local host. Any valid user name can be given.
<i>remuser</i>	Points to user names that are valid at the remote host. Any valid user name can be given.

<i>cmd</i>	Specifies the name of the command to be started at the remote host.
<i>fd2p</i>	Specifies an integer controlling the set up of communication channels. Integer options are as follows: <div> <div><b>Non-zero</b></div> <div>Indicates an auxiliary channel to a control process is set up, and the <i>fd2p</i> parameter points to the file descriptor for the channel. The control process provides diagnostic output from the remote command on this channel and also accepts bytes as signal numbers to be forwarded to the process group of the command.</div> <div><b>0</b></div> <div>Indicates the standard error (<b>stderr</b>) of the remote command is the same as standard output (<b>stdout</b>). No provision is made for sending arbitrary signals to the remote process. However, it is possible to send out-of-band data to the remote command.</div> </div>
<i>af</i>	The family argument is AF_INET, AF_INET6, or AF_UNSPEC. When either AF_INET or AF_INET6 is specified, this function will create a socket of the specified address family. When AF_UNSPEC is specified, it will try all possible address families until a connection can be established, and will return the associated socket of the connection.

## Return Values

Upon successful completion, the **rcmd\_af** subroutine returns a valid socket descriptor. Upon unsuccessful completion, the **rcmd\_af** subroutine returns a value of **-1**. The subroutine returns a **-1** if the effective user ID of the calling process is not the root user or if the subroutine is unsuccessful to resolve the host.

## Files

<i>/etc/services</i>	Contains the service names, ports, and socket type.
<i>/etc/hosts</i>	Contains host names and their addresses for hosts in a network.
<i>/etc/resolv.conf</i>	Contains the name server and domain name.

## Related Information

“rcmd Subroutine” on page 162, “rexec\_af Subroutine” on page 185, “rresvport\_af Subroutine” on page 187

---

## ReadFile Subroutine

### Purpose

Reads data from a socket.

### Syntax

```
#include <iocp.h>
boolean_t ReadFile (FileDescriptor, Buffer, ReadCount, AmountRead, Overlapped)
HANDLE FileDescriptor;
LPVOID Buffer;
DWORD ReadCount;
LPDWORD AmountRead;
LPOVERLAPPED Overlapped;
```

### Description

The **ReadFile** subroutine reads the number of bytes specified by the *ReadCount* parameter from the *FileDescriptor* parameter into the buffer indicated by the *Buffer* parameter. The number of bytes read is saved in the *AmountRead* parameter. The *Overlapped* parameter indicates whether or not the operation can be handled asynchronously.

The **ReadFile** subroutine returns a boolean (an integer) indicating whether or not the request has been completed.

The **ReadFile** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

**Note:** This subroutine only works to a socket file descriptor. It does not work with files or other file descriptors.

## Parameters

<i>FileDescriptor</i>	Specifies a valid file descriptor obtained from a call to the <b>socket</b> or <b>accept</b> subroutines.
<i>Buffer</i>	Specifies the buffer from which the data will be read.
<i>ReadCount</i>	Specifies the maximum number of bytes to read.
<i>AmountRead</i>	Specifies the number of bytes read. The parameter is set by the subroutine.
<i>Overlapped</i>	Specifies an overlapped structure indicating whether or not the request can be handled asynchronously.

## Return Values

Upon successful completion, the **ReadFile** subroutine returns a boolean indicating the request has been completed.

If the **ReadFile** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of 0 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable, see the link in the Related Information section of this document.

## Error Codes

The subroutine is unsuccessful if any of the following errors occur:

<b>EINPROGRESS</b>	The read request can not be immediately satisfied and will be handled asynchronously. A completion packet will be sent to the associated completion port upon completion.
<b>EAGAIN</b>	The read request cannot be immediately satisfied and cannot be handled asynchronously.
<b>EINVAL</b>	The <i>FileDescriptor</i> parameter is invalid.

## Examples

The following program fragment illustrates the use of the **ReadFile** subroutine to synchronously read data from a socket:

```
void buffer;  
int amount_read;  
b = ReadFile (34, &buffer, 128, &amount_read, NULL);
```

The following program fragment illustrates the use of the **ReadFile** subroutine to asynchronously read data from a socket:

```
void buffer;  
int amount_read;  
LPOVERLAPPED overlapped;  
b = ReadFile (34, &buffer, 128, &amount_read, overlapped);
```

**Note:** The request will only be handled asynchronously if it cannot be immediately satisfied.

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “CreateloCompletionPort Subroutine” on page 37, “WriteFile Subroutine” on page 258, “GetQueuedCompletionStatus Subroutine” on page 101, “GetMultipleCompletionStatus Subroutine” on page 83, and “PostQueuedCompletionStatus Subroutine” on page 161.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*

---

## recv Subroutine

### Purpose

Receives messages from connected sockets.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int recv (Socket,  
Buffer, Length, Flags)  
int Socket;  
void * Buffer;  
size_t Length;  
int Flags;
```

### Description

The **recv** subroutine receives messages from a connected socket. The **recvfrom** and **recvmsg** subroutines receive messages from both connected and unconnected sockets. However, they are usually used for unconnected sockets only.

The **recv** subroutine returns the length of the message. If a message is too long to fit in the supplied buffer, excess bytes may be truncated depending on the type of socket that issued the message.

If no messages are available at the socket, the **recv** subroutine waits for a message to arrive, unless the socket is nonblocking. If a socket is nonblocking, the system returns an error.

Use the **select** subroutine to determine when more data arrives.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

### Parameters

<i>Socket</i>	Specifies the socket descriptor.
<i>Buffer</i>	Specifies an address where the message should be placed.
<i>Length</i>	Specifies the size of the <i>Buffer</i> parameter.

*Flags* Points to a value controlling the message reception. The `/usr/include/sys/socket.h` file defines the *Flags* parameter. The argument to receive a call is formed by logically ORing one or more of the following values:

**MSG\_OOB**

Processes out-of-band data. The significance of out-of-band data is protocol-dependent.

**MSG\_PEEK**

Peeks at incoming data. The data continues to be treated as unread and will be read by the next call to **recv()** or a similar function.

**MSG\_WAITALL**

Requests that the function not return until the requested number of bytes have been read. The function can return fewer than the requested number of bytes only if a signal is caught, the connection is terminated, or an error is pending for the socket.

## Return Values

Upon successful completion, the **recv** subroutine returns the length of the message in bytes.

If the **recv** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Returns a 0 if the connection disconnects.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **recv** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	The remote peer forces the connection to be closed.
<b>EFAULT</b>	The data was directed to be received into a nonexistent or protected part of the process address space. The <i>Buffer</i> parameter is not valid.
<b>EINTR</b>	A signal interrupted the <b>recv</b> subroutine before any data was available.
<b>EINVAL</b>	The <b>MSG_OOB</b> flag is set and no out-of-band data is available.
<b>ENOBUF</b>	Insufficient resources are available in the system to perform the operation.
<b>ENOTCONN</b>	A receive is attempted on a <b>SOCK_STREAM</b> socket that is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	<b>MSG_OOB</b> flag is set for a <b>SOCK_DGRAM</b> socket, or <b>MSG_OOB</b> flag is set for any <b>AF_UNIX</b> socket.
<b>ETIMEDOUT</b>	The connection timed out during connection establishment, or there was a transmission timeout on an active connection.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **fgets** subroutine, **fputs** subroutine, **read** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **send** subroutine, **sendmsg** subroutine, **sendto** subroutine, **shutdown** subroutine, **socket** subroutine, **write** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## recvfrom Subroutine

### Purpose

Receives messages from sockets.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
ssize_t recvfrom  
(Socket, Buffer, Length, Flags, From, FromLength)  
int Socket;  
void * Buffer;  
size_t Length,  
int Flags;  
struct sockaddr * From;  
socklen_t * FromLength;
```

### Description

The **recvfrom** subroutine allows an application program to receive messages from unconnected sockets. The **recvfrom** subroutine is normally applied to unconnected sockets as it includes parameters that allow the calling program to specify the source point of the data to be received.

To return the source address of the message, specify a nonnull value for the *From* parameter. The *FromLength* parameter is a value-result parameter, initialized to the size of the buffer associated with the *From* parameter. On return, the **recvfrom** subroutine modifies the *FromLength* parameter to indicate the actual size of the stored address. The **recvfrom** subroutine returns the length of the message. If a message is too long to fit in the supplied buffer, excess bytes may be truncated depending on the type of socket that issued the message.

If no messages are available at the socket, the **recvfrom** subroutine waits for a message to arrive, unless the socket is nonblocking. If the socket is nonblocking, the system returns an error.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

### Parameters

<i>Socket</i>	Specifies the socket descriptor.
<i>Buffer</i>	Specifies an address where the message should be placed.
<i>Length</i>	Specifies the size of the <i>Buffer</i> parameter.

<i>Flags</i>	Points to a value controlling the message reception. The argument to receive a call is formed by logically ORing one or more of the values shown in the following list:
<b>MSG_OOB</b>	Processes out-of-band data. The significance of out-of-band data is protocol-dependent.
<b>MSG_PEEK</b>	Peeks at incoming data. The data continues to be treated as unread and will be read by the next call to <b>recv()</b> or a similar function.
<b>MSG_WAITALL</b>	Requests that the function not return until the requested number of bytes have been read. The function can return fewer than the requested number of bytes only if a signal is caught, the connection is terminated, or an error is pending for the socket.
<i>From</i>	Points to a socket structure, filled in with the source's address.
<i>FromLength</i>	Specifies the length of the sender's or source's address.

## Return Values

If the **recvfrom** subroutine is successful, the subroutine returns the length of the message in bytes.

If the call is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **recvfrom** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	The remote peer forces the connection to be closed.
<b>EFAULT</b>	The data was directed to be received into a nonexistent or protected part of the process address space. The buffer is not valid.
<b>EINTR</b>	The receive is interrupted by a signal delivery before any data is available.
<b>EINVAL</b>	The <b>MSG_OOB</b> flag is set but no out-of-band data is available.
<b>ENOBUF</b>	Insufficient resources are available in the system to perform the operation.
<b>ENOPROTOPT</b>	The protocol is not 64-bit supported.
<b>ENOTCONN</b>	A receive is attempted on a <b>SOCK_STREAM</b> socket that is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	<b>MSG_OOB</b> flag is set for a <b>SOCK_DGRAM</b> socket, or <b>MSG_OOB</b> flag is set for any <b>AF_UNIX</b> socket.
<b>ETIMEDOUT</b>	The connection timed out during connection establishment, or there was a transmission timeout on an active connection.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **fgets** subroutine, **fputs** subroutine, **read** subroutine, **recv** subroutine, **recvmsg** subroutine, **select** subroutine, **send** subroutine, **sendmsg** subroutine, **sendto** subroutine, **shutdown** subroutine, **socket** subroutine, **write** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## recvmsg Subroutine

### Purpose

Receives a message from any socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int recvmsg ( Socket, Message, Flags)
int Socket;
struct msghdr Message [ ];
int Flags;
```

### Description

The **recvmsg** subroutine receives messages from unconnected or connected sockets. The **recvmsg** subroutine returns the length of the message. If a message is too long to fit in the supplied buffer, excess bytes may be truncated depending on the type of socket that issued the message.

If no messages are available at the socket, the **recvmsg** subroutine waits for a message to arrive. If the socket is nonblocking and no messages are available, the **recvmsg** subroutine is unsuccessful.

Use the **select** subroutine to determine when more data arrives.

The **recvmsg** subroutine uses a **msghdr** structure to decrease the number of directly supplied parameters. The **msghdr** structure is defined in the **sys/socket.h** file. In BSD 4.3 Reno, the size and members of the **msghdr** structure have been modified. Applications wanting to start the old structure need to compile with **COMPAT\_43** defined. The default behavior is that of BSD 4.4.

All applications containing the **recvmsg** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Message</i>	Points to the address of the <b>msghdr</b> structure, which contains both the address for the incoming message and the space for the sender address.
<i>Flags</i>	Permits the subroutine to exercise control over the reception of messages. The <i>Flags</i> parameter used to receive a call is formed by logically ORing one or more of the values shown in the following list:

#### MSG\_OOB

Processes out-of-band data. The significance of out-of-band data is protocol-dependent.

#### MSG\_PEEK

Peeks at incoming data. The data continues to be treated as unread and will be read by the next call to **recv()** or a similar function.

#### MSG\_WAITALL

Requests that the function not return until the requested number of bytes have been read. The function can return fewer than the requested number of bytes only if a signal is caught, the connection is terminated, or an error is pending for the socket.

The **/sys/socket.h** file contains the possible values for the *Flags* parameter.

## Return Values

Upon successful completion, the length of the message in bytes is returned.

If the **recvmsg** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **recvmsg** subroutine is unsuccessful if any of the following error codes occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	The remote peer forces the connection to be closed.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINTR</b>	The <b>recvmsg</b> subroutine was interrupted by delivery of a signal before any data was available for the receive.
<b>EINVAL</b>	The length of the <b>msghdr</b> structure is invalid, or the <b>MSG_OOB</b> flag is set and no out-of-band data is available.
<b>EMSGSIZE</b>	The <i>msg_iovlen</i> member of the <b>msghdr</b> structure pointed to by <i>Message</i> is less than or equal to 0, or is greater than <b>IOV_MAX</b> .
<b>ENOBUF</b>	Insufficient resources are available in the system to perform the operation.
<b>ENOPROTOPT</b>	The protocol is not 64-bit supported.
<b>ENOTCONN</b>	A receive is attempted on a <b>SOCK_STREAM</b> socket that is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	<b>MSG_OOB</b> flag is set for a <b>SOCK_DGRAM</b> socket, or <b>MSG_OOB</b> flag is set for any <b>AF_UNIX</b> socket.
<b>ETIMEDOUT</b>	The connection timed out during connection establishment, or there was a transmission timeout on an active connection.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **no** command.

The **recv** subroutine, **recvfrom** subroutine, **select** subroutine, **send** subroutine, **sendmsg** subroutine, **sendto** subroutine, **shutdown** subroutine, **socket** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## res\_init Subroutine

### Purpose

Searches for a default domain name and Internet address.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
void res_init ( )
```

## Description

The **res\_init** subroutine reads the **/etc/resolv.conf** file for the default domain name and the Internet address of the initial hosts running the name server.

**Note:** If the **/etc/resolv.conf** file does not exist, the **res\_init** subroutine attempts name resolution using the local **/etc/hosts** file. If the system is not using a domain name server, the **/etc/resolv.conf** file should not exist. The **/etc/hosts** file should be present on the system even if the system is using a name server. In this instance, the file should contain the host IDs that the system requires to function even if the name server is not functioning.

The **res\_init** subroutine is one of a set of subroutines that form the resolver, a set of functions that translate domain names to Internet addresses. All resolver subroutines use the **/usr/include/resolv.h** file, which defines the **\_res** structure. The **res\_init** subroutine stores domain name information in the **\_res** structure. Three environment variables, **LOCALDOMAIN**, **RES\_TIMEOUT**, and **RES\_RETRY**, affect default values related to the **\_res** structure.

All applications containing the **res\_init** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

For more information on the **\_res** structure, see "Understanding Domain Name Resolution" in *AIX Version 6.1 Communications Programming Concepts*.

## Files

**/etc/resolv.conf**

Contains the name server and domain name.

**/etc/hosts**

Contains host names and their addresses for hosts in a network. This file is used to resolve a host name into an Internet address.

## Related Information

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_mkquery** subroutine, "res\_ninit Subroutine" on page 176, **res\_query** subroutine, **res\_search** subroutine, **res\_send** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## res\_mkquery Subroutine

### Purpose

Makes query messages for name servers.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
int res_mkquery (Operation, DomName, Class, Type, Data, DataLength)
int res_mkquery (Reserved, Buffer, BufferLength)
int Operation;
char * DomName;
int Class, Type;
char * Data;
int DataLength;
struct rrec * Reserved;
char * Buffer;
int BufferLength;
```

## Description

The **res\_mkquery** subroutine creates packets for name servers in the Internet domain. The subroutine also creates a standard query message. The *Buffer* parameter determines the location of this message.

The **res\_mkquery** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. Global information used by the resolver subroutines is kept in the **\_res** data structure. The **/usr/include/resolv.h** file contains the **\_res** structure definition.

All applications containing the **res\_mkquery** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Operation</i>	Specifies a query type. The usual type is <b>QUERY</b> , but the parameter can be set to any of the query types defined in the <b>arpa/nameser.h</b> file.
<i>DomName</i>	Points to the name of the domain. If the <i>DomName</i> parameter points to a single label and the <b>RES_DEFNAMES</b> structure is set, as it is by default, the subroutine appends the <i>DomName</i> parameter to the current domain name. The current domain name is defined by the name server in use or in the <b>/etc/resolv.conf</b> file.
<i>Class</i>	Specifies one of the following parameters:  <b>C_IN</b> Specifies the ARPA Internet.  <b>C_CHAOS</b> Specifies the Chaos network at MIT.

<i>Type</i>	Requires one of the following values:
	<b>T_A</b> Host address
	<b>T_NS</b> Authoritative server
	<b>T_MD</b> Mail destination
	<b>T_MF</b> Mail forwarder
	<b>T_CNAME</b> Canonical name
	<b>T_SOA</b> Start-of-authority zone
	<b>T_MB</b> Mailbox-domain name
	<b>T_MG</b> Mail-group member
	<b>T_MR</b> Mail-rename name
	<b>T_NULL</b> Null resource record
	<b>T_WKS</b> Well-known service
	<b>T_PTR</b> Domain name pointer
	<b>T_HINFO</b> Host information
	<b>T_MINFO</b> Mailbox information
	<b>T_MX</b> Mail-routing information
	<b>T_UINFO</b> User ( <b>finger</b> command) information
	<b>T_UID</b> User ID
	<b>T_GID</b> Group ID
<i>Data</i>	Points to the data that is sent to the name server as a search key. The data is stored as a character array.
<i>DataLength</i>	Defines the size of the array pointed to by the <i>Data</i> parameter.
<i>Reserved</i>	Specifies a reserved and currently unused parameter.
<i>Buffer</i>	Points to a location containing the query message.
<i>BufferLength</i>	Specifies the length of the message pointed to by the <i>Buffer</i> parameter.

## Return Values

Upon successful completion, the **res\_mkquery** subroutine returns the size of the query. If the query is larger than the value of the *BufferLength* parameter, the subroutine is unsuccessful and returns a value of -1.

## Files

<b>/etc/resolv.conf</b>	Contains the name server and domain name.
-------------------------	---

## Related Information

The **finger** command.

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, “res\_ninit Subroutine,” **res\_query** subroutine, **res\_search** subroutine, **res\_send** subroutine.

## Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

## res ninit Subroutine

## Purpose

Sets the default values for the members of the **res** structure.

## Library

## Standard C Library (**libc.a**)

## Syntax

```
#include <resolv.h>
```

```
int res_ninit (statp)
res state statp;
```

## Description

Reads the **/etc/resolv.conf** configuration file to get the default domain name, search list, and internet address of the local name server(s). It does this in order to re-initialize the resolver context for a given thread in a multi-threaded environment.

The **res\_ninit** subroutine sets the default values for the members of the **\_res** structure (defined in the **/usr/include/resolv.h** file) after reading the **/etc/resolv.conf** configuration file to get default domain name, search list, Internet address of the local name server(s), sort list, and options (for details, please refer to the **/etc/resolv.conf** file). If no name server is configured, the server address is set to **INADDR\_ANY** and the default domain name is obtained from the **gethostname** subroutine. It also allows the user to override **retrans**, **retry**, and local domain definition using three environment variables **RES\_TIMEOUT**, **RES\_RETRY**, and **LOCALDOMAIN**, respectively.

Using this subroutine, each thread can have unique local resolver context. Since the configuration file is read each time the subroutine is called, it is capable of tracking dynamic changes to the resolver state file. Changes include, addition or removal of the configuration file or any other modifications to this file and reflect the same for a given thread. The **res\_ninit** subroutine can also be used in single-threaded applications to detect dynamic changes to the resolver file even while the program is running (See the example section below). For more information on the **\_res** structure, see *Understanding Domain Name Resolution in AIX Version 6.1 Communications Programming Concepts*.

## Parameters

<i>statp</i>	Specifies the state to be initialized.
--------------	--

## Examples

```
# cat /etc/resolv.conf
domain in.ibm.com
nameserver 9.184.192.240
```

The following two examples use the **gethostbyname** system call to retrieve the host address of a system (florida.in.ibm.com) continuously. In the first example, **gethostbyname** is called (by a thread 'resolver') in a

multi-threaded environment. The second example is not. Before each call to **gethostbyname**, the **res\_ninit** subroutine is called to reflect dynamic changes to the configuration file.

```
1) #include <stdio.h>
    #include <netdb.h>
    #include <resolv.h>
    #include <pthread.h>

    void *resolver (void *arg);
    main( ) {
        pthread_t thid;
        if ( pthread_create(&thid, NULL, resolver, NULL) ) {
            printf("error in thread creation\n");
            exit( ); }
        pthread_exit(NULL);
    }

    void *resolver (void *arg) {
        struct hostent *hp;
        struct sockaddr_in client;
        while(1) {
            res_ninit(&_res);          /* res_init() with RES_INIT unset would NOT work here */

            hp = (struct hostent * ) gethostbyname("florida.in.ibm.com");
            bcopy(hp->h_addr_list[0], &client.sin_addr, sizeof(client.sin_addr));
            printf("hostname: %s\n", inet_ntoa(client.sin_addr));
        }
    }
```

If the **/etc/resolv.conf** file is present when the thread 'resolver' is invoked, the hostname will be resolved for that thread (using the nameserver 9.184.192.210) and the output will be hostname: 9.182.21.151.

If **/etc/resolv.conf** is not present, the output will be hostname: 0.0.0.0.

2) The changes to **/etc/resolv.conf** file are reflected even while the program is running

```
#include <stdio.h>
#include <resolv.h>
#include <sys.h>
#include <netdb.h>
#include <string.h>

main() {
    struct hostent *hp;
    struct sockaddr_in client;

    while (1) {
        res_ninit(&_res);

        hp = (struct hostent * ) gethostbyname("florida.in.ibm.com");
        bcopy(hp->h_addr_list[0], &client.sin_addr, sizeof(client.sin_addr));
        printf("hostname: %s\n", inet_ntoa(client.sin_addr));
    }
}
```

If **/etc/resolv.conf** is present while the program is running, the hostname will be resolved (using the nameserver 9.184.192.240) and the output will be hostname: 9.182.21.151.

If the **/etc/resolv.conf** file is not present, the output of the program will be hostname: 0.0.0.0.

**Note:** In the second example, the **res\_init** subroutine with **\_res.options = ~RES\_INIT** can be used instead of the **res\_ninit** subroutine.

## Files

The `/etc/resolv.conf` and `/etc/hosts` files.

## Related Information

The “dn\_comp Subroutine” on page 38, “dn\_expand Subroutine” on page 40, “\_getshort Subroutine” on page 26, “\_getlong Subroutine” on page 25, “\_putlong Subroutine” on page 27, “\_putshort Subroutine” on page 28, “res\_init Subroutine” on page 172, “res\_mkquery Subroutine” on page 173, “res\_query Subroutine,” “res\_search Subroutine” on page 180, “res\_send Subroutine” on page 182. Understanding Domain Name resolution

Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## res\_query Subroutine

### Purpose

Provides an interface to the server query mechanism.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int res_query (DomName, Class, Type, Answer, AnswerLength)
char * DomName;
int Class;
int Type;
u_char * Answer;
int AnswerLength;
```

### Description

The **res\_query** subroutine provides an interface to the server query mechanism. It constructs a query, sends it to the local server, awaits a response, and makes preliminary checks on the reply. The query requests information of the specified type and class for the fully-qualified domain name specified in the *DomName* parameter. The reply message is left in the answer buffer whose size is specified by the *AnswerLength* parameter, which is supplied by the caller.

The **res\_query** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. The **\_res** data structure contains global information used by the resolver subroutines. The `/usr/include/resolv.h` file contains the **\_res** structure definition.

All applications containing the **res\_query** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>DomName</i>	Points to the name of the domain. If the <i>DomName</i> parameter points to a single-component name and the <b>RES_DEFNAMES</b> structure is set, as it is by default, the subroutine appends the default domain name to the single-component name. The current domain name is defined by the name server in use or is specified in the <b>/etc/resolv.conf</b> file.
<i>Class</i>	Specifies one of the following values:  <b>C_IN</b> Specifies the ARPA Internet.  <b>C_CHAOS</b> Specifies the Chaos network at MIT.
<i>Type</i>	Requires one of the following values:  <b>T_A</b> Host address <b>T_NS</b> Authoritative server <b>T_MD</b> Mail destination <b>T_MF</b> Mail forwarder  <b>T_CNAME</b> Canonical name  <b>T_SOA</b> Start-of-authority zone <b>T_MB</b> Mailbox-domain name <b>T_MG</b> Mail-group member <b>T_MR</b> Mail-rename name  <b>T_NULL</b> Null resource record  <b>T_WKS</b> Well-known service  <b>T_PTR</b> Domain name pointer  <b>T_HINFO</b> Host information  <b>T_MINFO</b> Mailbox information  <b>T_MX</b> Mail-routing information  <b>T_UINFO</b> User ( <b>finger</b> command) information  <b>T_UID</b> User ID  <b>T_GID</b> Group ID
<i>Answer</i>	Points to an address where the response is stored.
<i>AnswerLength</i>	Specifies the size of the answer buffer.

## Return Values

Upon successful completion, the **res\_query** subroutine returns the size of the response. Upon unsuccessful completion, the **res\_query** subroutine returns a value of -1 and sets the **h\_errno** value to the appropriate error.

## Files

<b>/etc/resolv.conf</b>	Contains the name server and domain name.
-------------------------	---

## Related Information

The **finger** command.

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_search** subroutine, **res\_send** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## res\_search Subroutine

### Purpose

Makes a query and awaits a response.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int res_search (DomName, Class, Type, Answer, AnswerLength)
char * DomName;
int Class;
int Type;
u_char * Answer;
int AnswerLength;
```

### Description

The **res\_search** subroutine makes a query and awaits a response like the **res\_query** subroutine. However, it also implements the default and search rules controlled by the **RES\_DEFNAMES** and **RES\_DNSRCH** options.

The **res\_search** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. The **\_res** data structure contains global information used by the resolver subroutines. The **/usr/include/resolv.h** file contains the **\_res** structure definition.

All applications containing the **res\_search** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>DomName</i>	Points to the name of the domain. If the <i>DomName</i> parameter points to a single-component name and the <b>RES_DEFNAMES</b> structure is set, as it is by default, the subroutine appends the default domain name to the single-component name. The current domain name is defined by the name server in use or is specified in the <i>/etc/resolv.conf</i> file.
<i>Class</i>	<p>If the <b>RES_DNSRCH</b> bit is set, as it is by default, the <b>res_search</b> subroutine searches for host names in both the current domain and in parent domains.</p> <p>Specifies one of the following values:</p> <p><b>C_IN</b> Specifies the ARPA Internet.</p> <p><b>C_CHAOS</b> Specifies the Chaos network at MIT.</p>
<i>Type</i>	<p>Requires one of the following values:</p> <p><b>T_A</b> Host address</p> <p><b>T_NS</b> Authoritative server</p> <p><b>T_MD</b> Mail destination</p> <p><b>T_MF</b> Mail forwarder</p> <p><b>T_CNAME</b> Canonical name</p> <p><b>T_SOA</b> Start-of-authority zone</p> <p><b>T_MB</b> Mailbox-domain name</p> <p><b>T_MG</b> Mail-group member</p> <p><b>T_MR</b> Mail-rename name</p> <p><b>T_NULL</b> Null resource record</p> <p><b>T_WKS</b> Well-known service</p> <p><b>T_PTR</b> Domain name pointer</p> <p><b>T_HINFO</b> Host information</p> <p><b>T_MINFO</b> Mailbox information</p> <p><b>T_MX</b> Mail-routing information</p> <p><b>T_UINFO</b> User (<b>finger</b> command) information</p> <p><b>T_UID</b> User ID</p> <p><b>T_GID</b> Group ID</p>
<i>Answer</i>	Points to an address where the response is stored.
<i>AnswerLength</i>	Specifies the size of the answer buffer.

## Return Values

Upon successful completion, the **res\_search** subroutine returns the size of the response. Upon unsuccessful completion, the **res\_search** subroutine returns a value of -1 and sets the **h\_errno** value to the appropriate error.

## Files

**/etc/resolv.conf**

Contains the name server and domain name.

## Related Information

The **finger** command.

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_query** subroutine, **res\_send** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## res\_send Subroutine

### Purpose

Sends a query to a name server and retrieves a response.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int res_send (MessagePtr, MessageLength, Answer, AnswerLength)
char * MsgPtr;
int MsgLength;
char * Answer;
int AnswerLength;
```

### Description

The **res\_send** subroutine sends a query to name servers and calls the **res\_init** subroutine if the **RES\_INIT** option of the **\_res** structure is not set. This subroutine sends the query to the local name server and handles time outs and retries.

The **res\_send** subroutine is one of a set of subroutines that form the resolver, a set of functions that resolve domain names. Global information used by the resolver subroutines is kept in the **\_res** structure. The **/usr/include/resolv.h** file contains the **\_res** structure definition.

All applications containing the **res\_send** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>MessagePtr</i>	Points to the beginning of a message.
<i>MessageLength</i>	Specifies the length of the message.

<i>Answer</i>	Points to an address where the response is stored.
<i>AnswerLength</i>	Specifies the size of the answer area.

## Return Values

Upon successful completion, the **res\_send** subroutine returns the length of the message.

If the **res\_send** subroutine is unsuccessful, the subroutine returns a -1.

## Files

<b>/etc/resolv.conf</b>	Contains general name server and domain name information.
-------------------------	---

## Related Information

The **dn\_comp** subroutine, **dn\_expand** subroutine, **\_getlong** subroutine, **\_getshort** subroutine, **putlong** subroutine, **putshort** subroutine, **res\_init** subroutine, **res\_mkquery** subroutine, “res\_ninit Subroutine” on page 176, **res\_query** subroutine, **res\_search** subroutine.

Sockets Overview and Understanding Domain Name Resolution in *AIX Version 6.1 Communications Programming Concepts*.

---

## rexec Subroutine

### Purpose

Allows command execution on a remote host.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int rexec ( Host, Port, User, Passwd, Command, ErrFileDescParam)
char **Host;
int Port;
char *User, *Passwd,
*Command;
int *ErrFileDescParam;
```

### Description

The **rexec** subroutine allows the calling process to start commands on a remote host.

If the **rexec** connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process and is given to the remote command as standard input and standard output.

All applications containing the **rexec** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Host</i>	Contains the name of a remote host that is listed in the <b>/etc/hosts</b> file or <b>/etc/resolv.conf</b> file. If the name of the host is not found in either file, the <b>rexec</b> subroutine is unsuccessful.
<i>Port</i>	Specifies the well-known DARPA Internet port to use for the connection. A pointer to the structure that contains the necessary port can be obtained by issuing the following library call:  <code>getservbyname("exec","tcp")</code>
<i>User and Passwd</i>	Points to a user ID and password valid at the host. If these parameters are not supplied, the <b>rexec</b> subroutine takes the following actions until finding a user ID and password to send to the remote host: <ol style="list-style-type: none"><li>1. Searches the current environment for the user ID and password on the remote host.</li><li>2. Searches the user's home directory for a file called <b>\$HOME/.netrc</b> that contains a user ID and password.</li><li>3. Prompts the user for a user ID and password.</li></ol>
<i>Command</i>	Points to the name of the command to be executed at the remote host.
<i>ErrFileDescParam</i>	Specifies one of the following values:  <b>Non-zero</b> Indicates an auxiliary channel to a control process is set up, and a descriptor for it is placed in the <i>ErrFileDescParam</i> parameter. The control process provides diagnostic output from the remote command on this channel and also accepts bytes as signal numbers to be forwarded to the process group of the command. This diagnostic information does not include remote authorization failure, since this connection is set up after authorization has been verified.  <b>0</b> Indicates the standard error of the remote command is the same as standard output, and no provision is made for sending arbitrary signals to the remote process. In this case, however, it may be possible to send out-of-band data to the remote command.

## Return Values

Upon successful completion, the system returns a socket to the remote command.

If the **rexec** subroutine is unsuccessful, the system returns a -1 indicating that the specified host name does not exist.

## Files

<b>/etc/hosts</b>	Contains host names and their addresses for hosts in a network. This file is used to resolve a host name into an Internet address.
<b>/etc/resolv.conf</b>	Contains the name server and domain name.
<b>\$HOME/.netrc</b>	Contains automatic login information.

## Related Information

The **getservbyname** subroutine, **rcmd** subroutine, **rresvport** subroutine, **ruserok** subroutine.

The **rexecd** daemon.

Transmission Control Protocol/Internet Protocol in *Networks and communication management*.

## rexec\_af Subroutine

### Purpose

Allows command execution on a remote host.

### Syntax

```
int rexec_af(char **ahost, unsigned short rport, const char *name,  
             const char *pass, const char *cmd, int *fd2p, int af)
```

### Description

The **rexec\_af** subroutine allows the calling process to start commands on a remote host. It behaves the same as the existing **rexec()** function, but instead of creating only an AF\_INET TCP socket, it can also create an AF\_INET6 TCP socket.

The **rexec\_af** subroutine is useful because the existing **rexec()** function cannot transparently use AF\_INET6 sockets. This is because an application would not be prepared to handle AF\_INET6 addresses returned by functions such as **getpeername()** on the file descriptor created by **rexec()**.

If the **rexec\_af** connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process and is given to the remote command as standard input and standard output.

All applications containing the **rexec\_af** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>ahost</i>	Contains the name of a remote host that is listed in the <b>/etc/hosts</b> file or <b>/etc/resolv.config</b> file. If the name of the host is not found in either file, the <b>rexec</b> subroutine is unsuccessful.
<i>rport</i>	Specifies the well-known DARPA Internet port to use for the connection. A pointer to the structure that contains the necessary port can be obtained by issuing the following library call: <code>getservbyname("exec", "tcp")</code>
<i>name and pass</i>	Points to a valid user ID and password at the host. If these parameters are not supplied, the <b>rexec_af</b> subroutine takes the following actions until it finds a user ID and password to send to the remote host: <ol style="list-style-type: none"><li>1. Searches the current environment for the user ID and password on the remote host.</li><li>2. Searches the user's home directory for a file called <b>\$HOME/.netrc</b> that contains a user ID and password.</li><li>3. Prompts the user for a user ID and password.</li></ol>
<i>cmd</i>	Points to the name of the command to be executed at the remote host.

<i>fd2p</i>	Specifies one of the following values:
	<b>Non-zero</b> Indicates that an auxiliary channel to a control process is set up, and a descriptor for it is placed in the <i>fd2p</i> parameter. The control process provides diagnostic output from the remote command on this channel and also accepts bytes as signal numbers to be forwarded to the process group of the command. This diagnostic information does not include remote authorization failure, since this connection is set up after authorization has been verified.
	<b>0</b> Indicates that the standard error of the remote command is the same as standard output, and no provision is made for sending arbitrary signals to the remote process. In this case, however, it might be possible to send out-of-band data to the remote command.
<i>af</i>	The family argument is AF_INET, AF_INET6, or AF_UNSPEC. When either AF_INET or AF_INET6 is specified, this subroutine will create a socket of the specified address family. When AF_UNSPEC is specified, it will try all possible address families until a connection can be established, and will return the associated socket of the connection.

## Return Values

Upon successful completion, the system returns a socket to the remote command. If the **rexec\_af** subroutine is unsuccessful, the system returns a -1, indicating that the specified host name does not exist.

## Files

<i>/etc/hosts</i>	Contains host names and their addresses for hosts in a network. This file is used to resolve a host name into an Internet address.
<i>/etc/resolv.conf</i>	Contains the name server and domain name.
<i>\$HOME/.netrc</i>	Contains automatic login information.

## Related Information

“rcmd\_af Subroutine” on page 164, “rexec Subroutine” on page 183, “rresvport\_af Subroutine” on page 187.

---

## rresvport Subroutine

### Purpose

Retrieves a socket with a privileged address.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int rresvport ( Port)
int *Port;
```

### Description

The **rresvport** subroutine obtains a socket with a privileged address bound to the socket. A privileged Internet port is one that falls in a range between 0 and 1023.

Only processes with an effective user ID of root user can use the **rresvport** subroutine. An authentication scheme based on remote port numbers is used to verify permissions.

If the connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process.

All applications containing the **rresvport** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*Port* Specifies the port to use for the connection.

## Return Values

Upon successful completion, the **rresvport** subroutine returns a valid, bound socket descriptor.

If the **rresvport** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **rresvport** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EAGAIN</b>	All network ports are in use.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EMFILE</b>	Two hundred file descriptors are currently open.
<b>ENFILE</b>	The system file table is full.
<b>ENOBUFS</b>	Insufficient buffers are available in the system to complete the subroutine.

## Files

**/etc/services** Contains the service names.

## Related Information

The **rcmd** subroutine, **ruserok** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## rresvport\_af Subroutine

### Purpose

Retrieves a socket with a privileged address.

### Syntax

```
int rresvport_af(int *port, int family);
```

### Description

The **rresvport\_af** subroutine obtains a socket with a privileged address bound to the socket. A privileged Internet port is one that falls in a range between 0 and 1023.

This subroutine is similar to the existing `rresvport()` subroutine, except that `rresvport_af` also takes and address family as an argument. This function is capable of creating either an `AF_INET/TCP` or an `AF_INET6/TCP` socket.

Only processes with an effective user ID of root user can use the **rresvport** subroutine. An authentication scheme based on remote port numbers is used to verify permissions.

If the connection succeeds, a socket in the Internet domain of type **SOCK\_STREAM** is returned to the calling process.

All applications containing the **rresvport** subroutine must be compiled with the `_BSD` macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD `libbsd.a` library.

## Parameters

<i>port</i>	Specifies the port to use for the connection.
<i>family</i>	Specifies either <code>AF_INET</code> or <code>AF_INET6</code> to accommodate the appropriate version.

## Return Values

Upon successful completion, the **rresvport\_af** subroutine returns a valid, bound socket descriptor.

If the **rresvport\_af** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

<b>EAFNOSUPPORT</b>	The address family is not supported.
<b>EAGAIN</b>	All network ports are in use.
<b>EMFILE</b>	Two hundred file descriptors are currently open.
<b>ENFILE</b>	The system file table is full.
<b>ENOBUFS</b>	Insufficient buffers are available in the system to complete the subroutine.

## Files

<code>/etc/services</code>	Contains the service names.
----------------------------	-----------------------------

## Related Information

“`rcmd_af` Subroutine” on page 164, “`rexec_af` Subroutine” on page 185, “`rresvport` Subroutine” on page 186.

---

## ruserok Subroutine

### Purpose

Allows servers to authenticate clients.

### Library

Standard C Library (**libc.a**)

## Syntax

```
int ruserok (Host, RootUser, RemoteUser, LocalUser)
char * Host;
int RootUser;
char * RemoteUser,
* LocalUser;
```

## Description

The **ruserok** subroutine allows servers to authenticate clients requesting services.

Always specify the host name. If the local domain and remote domain are the same, specifying the domain parts is optional. To determine the domain of the host, use the **gethostname** subroutine.

All applications containing the **ruserok** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Host</i>	Specifies the name of a remote host. The <b>ruserok</b> subroutine checks for this host in the <b>/etc/host.equiv</b> file. Then, if necessary, the subroutine checks a file in the user's home directory at the server called <b>/\$HOME/.rhosts</b> for a host and remote user ID.
<i>RootUser</i>	Specifies a value to indicate whether the effective user ID of the calling process is a root user. A value of 0 indicates the process does not have a root user ID. A value of 1 indicates that the process has local root user privileges, and the <b>/etc/hosts.equiv</b> file is not checked.
<i>RemoteUser</i>	Points to a user name that is valid at the remote host. Any valid user name can be specified.
<i>LocalUser</i>	Points to a user name that is valid at the local host. Any valid user name can be specified.

## Return Values

The **ruserok** subroutine returns a 0, if the subroutine successfully locates the name specified by the *Host* parameter in the **/etc/hosts.equiv** file or the IDs specified by the *Host* and *RemoteUser* parameters are found in the **/\$HOME/.rhosts** file.

If the name specified by the *Host* parameter was not found, the **ruserok** subroutine returns a -1.

## Files

<b>/etc/services</b>	Contains service names.
<b>/etc/host.equiv</b>	Specifies foreign host names.
<b>/\$HOME/.rhosts</b>	Specifies the remote users of a local user account.

## Related Information

The **rlogind** command, **rshd** command.

The **gethostname** subroutine, **rcmd** subroutine, **rresvport** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## sctp\_opt\_info Subroutine

### Purpose

Passes information both into and out of SCTP stack.

### Library

Standard C Library (**libc.a**)

### Syntax

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

int sctp_opt_info(sd, id, opt, *arg_size, *size);
int sd;
sctp_assoc_t id;
int opt;
void *arg_size;
size_t *size;
```

### Description

Applications use the **sctp\_opt\_info** subroutine to get information about various SCTP socket options from the stack. For the sockets with multiple associations, the association ID can be specified to apply the operation on any particular association of a socket. Because an SCTP association supports multihoming, this operation can be used to specify any particular peer address using a **sockaddr\_storage** structure. In this case, the result of the operation will be applied to only that particular peer address.

### Implementation Specifics

The **sctp\_opt\_info** subroutine is part of Base Operating System (BOS) Runtime.

### Parameters

<i>sd</i>	Specifies the UDP style socket descriptor returned from the <b>socket</b> system call.
<i>id</i>	Specifies the identifier of the association to query.
<i>opt</i>	Specifies the socket option to get.
<i>arg_size</i>	Specifies an option specific structure buffer provided by the caller.
<i>size</i>	Specifies the size of the option returned.

### Return Values

Upon successful completion, the **sctp\_opt\_info** subroutine returns 0.

If the **sctp\_opt\_info** subroutine is unsuccessful, the subroutine handler returns a value of -1 to the calling program and sets **errno** to the appropriate error code.

### Error Codes

The **sctp\_opt\_info** subroutine is unsuccessful if any of the following errors occurs:

<b>EFAULT</b>	Indicates that the user has insufficient authority to access the data, or the address specified in the <i>uaddr</i> parameter is not valid.
<b>EIO</b>	Indicates that a permanent I/O error occurred while referencing data.
<b>ENOMEM</b>	Indicates insufficient memory for the required paging operation.

<b>ENOSPC</b>	Indicates insufficient file system or paging space.
<b>ENOBUFS</b>	Insufficient resources were available in the system to complete the call.
<b>ENOPROTOPT</b>	Protocol not available.
<b>ENOTSOCK</b>	Indicates that the user has tried to do a socket operation on a non-socket.

## Related Information

The “sctpctrl Subroutine” on page 192, “sctp\_peeloff Subroutine.”

Stream Control Transmission Protocol in *Networks and communication management*.

---

## sctp\_peeloff Subroutine

### Purpose

Branches off an association into a separate socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

int sctp_peeloff(sd, *assoc_id);
int sd;
sctp_assoc_t *assoc_id;
```

### Description

An application uses the **sctp\_peeloff** subroutine when it wants to branch-off an existing association into a separate socket/file descriptor. It returns a new socket descriptor, which in turn can be used to send and receive subsequent SCTP packets. After it has been branched off, an association becomes completely independent of the original socket. Any subsequent data or control operations to that association must be passed using the new socket descriptor. Also, a close on the original socket descriptor will not close the new socket descriptor branched out of the association.

All the associations under the same socket share the same socket buffer space of the socket that they belong to. If an association gets branched off to a new socket using **sctp\_peeloff**, then it inherits the socket buffer space associated with the new socket descriptor. This way, the association that got peeled off keeps more buffer space.

### Implementation Specifics

The **sctp\_peeloff** subroutine is part of Base Operating System (BOS) Runtime.

### Parameters

<i>sd</i>	Specifies the UDP style socket descriptor returned from the <b>socket</b> system call.
<i>assoc_id</i>	Specifies the identifier of the association that is to be branched-off to a separate socket descriptor.

## Return Values

Upon successful completion, the **sctp\_peeloff** subroutine returns the nonnegative socket descriptor of the branched-off socket.

If the **sctp\_peeloff** subroutine is unsuccessful, the subroutine handler returns a value of -1 to the calling program and moves an error code to the **errno** global variable.

## Error Codes

The **sctp\_peeloff** subroutine is unsuccessful if any of the following errors occurs:

<b>EINVAL</b>	Invalid argument.
<b>EBADF</b>	Bad file descriptor.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>ESOCKTNOSUPPORT</b>	The socket in the specified address family is not supported.
<b>EMFILE</b>	The per-process descriptor table is full.
<b>ENOBUFS</b>	Insufficient resources were available in the system to complete the call.
<b>ECONNABORTED</b>	The client aborted the connection.

## Related Information

“sctpctrl Subroutine,” “sctp\_opt\_info Subroutine” on page 190

Stream Control Transmission Protocol in *Networks and communication management*.

---

## sctpctrl Subroutine

### Purpose

Controls and configures SCTP.

### Syntax

```
sctpctrl {load|unload|set}  
sctpctrl stats [reset] [interval]  
  
sctpctrl set {name=value|default [name]}  
  
sctpctrl get [name]
```

### Description

The **sctpctrl** subroutine controls and configures the SCTP kernel extension. This subroutine can be used to load and unload the SCTP kernel extension. It can also be used to dump SCTP data, and set and retrieve various SCTP tunables. In addition, the **sctpctrl** subroutine can be used to read and reset the SCTP specific network statistics.

### Parameters

<i>load</i>	Loads the SCTP kernel extension if not loaded.
<i>unload</i>	Unloads the SCTP kernel extension if loaded.
<i>stats</i> [ <b>reset</b> ] [ <i>interval</i> ]	Displays SCTP statistics. The optional <b>reset</b> command will clear (0) the statistics. If the <i>interval</i> parameter (in seconds) is added, the program does not exit; instead, it outputs the statistics every <i>interval</i> seconds.

**set** {*name=value*|*default* [*name*]} Sets the SCTP tunable to a value. If *default* is specified, all the tunables are set to their default values. If optional [*name*] is specified followed by *default*, the tunable described by *name* is set to its default value.

**get** [*name*] Gets the value of the tunable described by its optional name parameter. If the *name* parameter is not specified, *get* gets the values of all the tunables.

## Examples

1. To load the **sctp** kernel extension, enter:

```
sctpctrl load
```

2. To unload the **sctp** kernel extension, enter:

```
sctpctrl unload
```

3. To reset the SCTP statistics, enter:

```
sctpctrl stats reset
```

This command will zero-out all the SCTP statistics.

4. To get the values of the SCTP tunable, enter:

```
sctpctrl get
```

This will list all the SCTP tunables and their values. Here is a sample output.

```
sctp_assoc_maxerr = 10
sctp_cookie_life = 60
sctp_delack_timer = 4
sctp_dontdelayack = 1
sctp_ecn = 1
sctp_ephemeral_high = 65535
sctp_ephemeral_low = 32768
sctp_instreams = 2048
sctp_maxburst = 8
sctp_outstreams = 10
sctp_path_maxerr = 5
sctp_pmtu_discover = 1
sctp_rttmax = 60
sctp_rttmin = 1
sctp_recvspace = 65536
sctp_sendspace = 65536
sctp_send_fewsacks = 0
```

5. To set **sctp\_path\_maxerr** to a value of 6, enter:

```
sctpctrl set sctp_path_maxerr=6
```

## Files

<b>/usr/sbin/sctpctrl</b>	Contains the <b>sctpctrl</b> command.
<b>/usr/lib/drivers/sctp</b>	Contains the SCTP kernel extension.

## Related Information

The “sctp\_peeloff Subroutine” on page 191, “sctp\_opt\_info Subroutine” on page 190.

Stream Control Transmission Protocol in *Networks and communication management*.

---

## send Subroutine

### Purpose

Sends messages from a connected socket.

## Library

Standard C Library (**libc.a**)

## Syntax

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

```
int send (Socket,
Message, Length, Flags)
int Socket;
const void * Message;
size_t Length;
int Flags;
```

## Description

The **send** subroutine sends a message only when the socket is connected. This subroutine on a socket is not thread safe. The **sendto** and **sendmsg** subroutines can be used with unconnected or connected sockets.

To broadcast on a socket, first issue a **setsockopt** subroutine using the **SO\_BROADCAST** option to gain broadcast permissions.

Specify the length of the message with the *Length* parameter. If the message is too long to pass through the underlying protocol, the system returns an error and does not transmit the message.

No indication of failure to deliver is implied in a **send** subroutine. A return value of -1 indicates some locally detected errors.

If no space for messages is available at the sending socket to hold the message to be transmitted, the **send** subroutine blocks unless the socket is in a nonblocking I/O mode. Use the **select** subroutine to determine when it is possible to send more data.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name for the socket.
<i>Message</i>	Points to the address of the message to send.
<i>Length</i>	Specifies the length of the message in bytes.
<i>Flags</i>	Allows the sender to control the transmission of the message. The <i>Flags</i> parameter used to send a call is formed by logically ORing one or both of the values shown in the following list:

### **MSG\_OOB**

Processes out-of-band data on sockets that support **SOCK\_STREAM** communication.

### **MSG\_DONTROUTE**

Sends without using routing tables.

### **MSG\_MPEG2**

Indicates that this block is a MPEG2 block. This flag is valid **SOCK\_CONN\_DGRAM** types of sockets only.

## Return Values

Upon successful completion, the **send** subroutine returns the number of characters sent.

If the **send** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EACCES</b>	Write access to the named socket is denied, or the socket trying to send a broadcast packet does not have broadcast capability.
<b>EADDRNOTAVAIL</b>	The specified address is not a valid address.
<b>EAFNOSUPPORT</b>	The specified address is not a valid address for the address family of this socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	A connection was forcibly closed by a peer.
<b>EDESTADDRREQ</b>	The socket is not in connection-mode and no peer address is set.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EHOSTUNREACH</b>	The destination host cannot be reached.
<b>EINTR</b>	A signal interrupted <b>send</b> before any data was transmitted.
<b>EINVAL</b>	The <i>Length</i> parameter is invalid.
<b>EISCONN</b>	A <b>SOCK_DGRAM</b> socket is already connected.
<b>EMSGSIZE</b>	The message is too large to be sent all at once, as the socket requires.
<b>ENETUNREACH</b>	The destination network is not reachable.
<b>ENOBUFS</b>	Insufficient resources were available in the system to perform the operation.
<b>ENOENT</b>	The path name does not name an existing file, or the path name is an empty string.
<b>ENOMEM</b>	The available data space in memory is not large enough to hold group/ACL information.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The socket argument is associated with a socket that does not support one or more of the values set in <i>Flags</i> .
<b>EPIPE</b>	An attempt was made to send on a socket that was connected, but the connection has been shut down either by the remote peer or by this side of the connection. If the socket is of type <b>SOCK_STREAM</b> , the <b>SIGPIPE</b> signal is generated to the calling process.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **connect** subroutine, **getsockopt** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **sendmsg** subroutine, **sendto** subroutine, **setsockopt** subroutine, **shutdown** subroutine, **socket** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## sendmsg Subroutine

### Purpose

Sends a message from a socket using a message structure.

### Library

Standard C Library (**libc.a**)

## Syntax

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

```
int sendmsg ( Socket, Message, Flags)
int Socket;
const struct msghdr Message [ ];
int Flags;
```

## Description

The **sendmsg** subroutine sends messages through connected or unconnected sockets using the **msghdr** message structure. The **/usr/include/sys/socket.h** file contains the **msghdr** structure and defines the structure members. In BSD 4.4, the size and members of the **msghdr** message structure have been modified. Applications wanting to start the old structure need to compile with **COMPAT\_43** defined. The default behaviour is that of BSD 4.4.

To broadcast on a socket, the application program must first issue a **setsockopt** subroutine using the **SO\_BROADCAST** option to gain broadcast permissions.

The **sendmsg** subroutine supports only 15 message elements.

All applications containing the **sendmsg** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

The **sendmsg** routine supports IPv6 ancillary data elements as defined in the Advanced Sockets API for IPv6.

## Parameters

<i>Socket</i>	Specifies the socket descriptor.
<i>Message</i>	Points to the <b>msghdr</b> message structure containing the message to be sent.
<i>Flags</i>	Allows the sender to control the message transmission. The <b>sys/socket.h</b> file contains the <i>Flags</i> parameter. The <i>Flags</i> parameter used to send a call is formed by logically ORing one or both of the following values:

### MSG\_OOB

Processes out-of-band data on sockets that support **SOCK\_STREAM**.

**Note:** The following value is not for general use. It is an administrative tool used for debugging or for routing programs.

### MSG\_DONTROUTE

Sends without using routing tables.

### MSG\_MPEG2

Indicates that this block is a MPEG2 block. It only applies to **SOCK\_CONN\_DGRAM** types of sockets only.

## Return Values

Upon successful completion, the **sendmsg** subroutine returns the number of characters sent.

If the **sendmsg** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

The **sendmsg** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EACCES</b>	Write access to the named socket is denied, or the socket trying to send a broadcast packet does not have broadcast capability.
<b>EADDRNOTAVAIL</b>	The specified address is not a valid address.
<b>EAFNOSUPPORT</b>	The specified address is not a valid address for the address family of this socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	A connection was forcibly closed by a peer.
<b>EDESTADDRREQ</b>	The socket is not in connection-mode and does not have its peer address set.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EHOSTUNREACH</b>	The destination host cannot be reached.
<b>EINTR</b>	A signal interrupted <b>sendmsg</b> before any data was transmitted.
<b>EINVAL</b>	The length of the <b>msghdr</b> structure is invalid.
<b>EISCONN</b>	A <b>SOCK_DGRAM</b> socket is already connected.
<b>EMSGSIZE</b>	The message is too large to be sent all at once (as the socket requires), or the <b>msg_iovlen</b> member of the <b>msghdr</b> structure pointed to by <i>Message</i> is less than or equal to 0 or is greater than <b>IOV_MAX</b> .
<b>ENOENT</b>	The path name does not name an existing file, or the path name is an empty string.
<b>ENETUNREACH</b>	The destination network is not reachable.
<b>ENOBUFS</b>	The system ran out of memory for an internal data structure.
<b>ENOMEM</b>	The available data space in memory is not large enough to hold group/ACL information.
<b>ENOPROTOOPT</b>	The protocol is not 64-bit supported.
<b>ENOTCONN</b>	The socket is in connection-mode but is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The socket argument is associated with a socket that does not support one or more of the values set in flags.
<b>EPIPE</b>	An attempt was made to send on a socket that was connected, but the connection has been shut down either by the remote peer or by this side of the connection. If the socket is of type <b>SOCK_STREAM</b> , the <b>SIGPIPE</b> signal is generated to the calling process.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **no** command.

The **connect** subroutine, **getsockopt** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **send** subroutine, **sendto** subroutine, **setsockopt** subroutine, **shutdown** subroutine, **socket** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## sendto Subroutine

### Purpose

Sends messages through a socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```

int sendto
(Socket, Message, Length,
Flags, To, ToLength)
int Socket;
const void * Message;
size_t Length;
int Flags;
const struct sockaddr * To;
socklen_t ToLength;

```

## Description

The **sendto** subroutine allows an application program to send messages through an unconnected socket by specifying a destination address.

To broadcast on a socket, first issue a **setsockopt** subroutine using the **SO\_BROADCAST** option to gain broadcast permissions.

Provide the address of the target using the *To* parameter. Specify the length of the message with the *Length* parameter. If the message is too long to pass through the underlying protocol, the error **EMSGSIZE** is returned and the message is not transmitted.

If the **sending** socket has no space to hold the message to be transmitted, the **sendto** subroutine blocks the message unless the socket is in a nonblocking I/O mode.

Use the **select** subroutine to determine when it is possible to send more data.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name for the socket.
<i>Message</i>	Specifies the address containing the message to be sent.
<i>Length</i>	Specifies the size of the message in bytes.
<i>Flags</i>	Allows the sender to control the message transmission. The <i>Flags</i> parameter used to send a call is formed by logically ORing one or both of the following values:

### MSG\_OOB

Processes out-of-band data on sockets that support **SOCK\_STREAM**.

### Note:

### MSG\_DONTROUTE

Sends without using routing tables.

<i>To</i>	The <b>/usr/include/sys/socket.h</b> file defines the <i>Flags</i> parameter. Specifies the destination address for the message. The destination address is a <b>sockaddr</b> structure defined in the <b>/usr/include/sys/socket.h</b> file.
<i>ToLength</i>	Specifies the size of the destination address.

## Return Values

Upon successful completion, the **sendto** subroutine returns the number of characters sent.

If the **sendto** subroutine is unsuccessful, the system returns a value of -1, and the **errno** global variable is set to indicate the error.

## Error Codes

The subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EACCES</b>	Write access to the named socket is denied, or the socket trying to send a broadcast packet does not have broadcast capability.
<b>EADDRNOTAVAIL</b>	The specified address is not a valid address.
<b>EAfnOSUPPORT</b>	The specified address is not a valid address for the address family of this socket.
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ECONNRESET</b>	A connection was forcibly closed by a peer.
<b>EDESTADDRREQ</b>	The socket is not in connection-mode and no peer address is set.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EHOSTUNREACH</b>	The destination host cannot be reached.
<b>EINTR</b>	A signal interrupted <b>sendto</b> before any data was transmitted.
<b>EINVAL</b>	The <i>Length</i> or <i>ToLength</i> parameter is invalid.
<b>EISCONN</b>	A <b>SOCK_DGRAM</b> socket is already connected.
<b>EMSGSIZE</b>	The message is too large to be sent all at once as the socket requires.
<b>ENETUNREACH</b>	The destination network is not reachable.
<b>ENOBUFS</b>	The system ran out of memory for an internal data structure.
<b>ENOENT</b>	The path name does not name an existing file, or the path name is an empty string.
<b>ENOMEM</b>	The available data space in memory is not large enough to hold group/ACL information.
<b>ENOPROTOOPT</b>	The protocol is not 64-bit supported.
<b>ENOTCONN</b>	The socket is in connection-mode but is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EOPNOTSUPP</b>	The socket argument is associated with a socket that does not support one or more of the values set in <i>Flags</i> .
<b>EPIPE</b>	An attempt was made to send on a socket that was connected, but the connection has been shut down either by the remote peer or by this side of the connection. If the socket is of type <b>SOCK_STREAM</b> , the <b>SIGPIPE</b> signal is generated to the calling process.
<b>EWOULDBLOCK</b>	The socket is marked nonblocking, and no connections are present to be accepted.

## Related Information

The **getsockopt** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **send** subroutine, **sendmsg** subroutine, **setsockopt** subroutine, **shutdown** subroutine, **socket** subroutine.

Sending UNIX Datagrams Example Program, Sending Internet Datagrams Example Program, Sockets Overview, Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## send\_file Subroutine

### Purpose

Sends the contents of a file through a socket.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
ssize_t send_file(Socket_p, sf_iobuf, flags)
```

```
int * Socket_p;  
struct sf_parms * sf_iobuf;  
uint_t flags;
```

## Description

The **send\_file** subroutine sends data from the opened file specified in the *sf\_iobuf* parameter, over the connected socket pointed to by the *Socket\_p* parameter.

**Note:** Currently, the **send\_file** only supports the TCP/IP protocol (SOCK\_STREAM socket in AF\_INET). An error will be returned when this function is used on any other types of sockets.

## Parameters

*Socket\_p*      Points to the socket descriptor of the socket which the file will be sent to.  
**Note:** This is different from most of the socket functions.

*sf\_iobuf* Points to a *sf\_parms* structure defined as follows:

```
/*
 * Structure for the send_file system call
 */
#ifdef _64BIT_
#define SF_INT64(x)      int64_t x;
#define SF_UINT64(x)     uint64_t x;
#else
#ifdef _LONG_LONG
#define SF_INT64(x)      int64_t x;
#define SF_UINT64(x)     uint64_t x;
#else
#define SF_INT64(x)      int filler_##x; int x;
#define SF_UINT64(x)     int filler_##x; uint_t x;
#endif
#endif

struct sf_parms {
    /* ----- header parms ----- */
    void      *header_data;          /* Input/Output. Points to header buf */
    uint_t     header_length;        /* Input/Output. Length of the header */
    /* ----- file parms ----- */
    int        file_descriptor;      /* Input. File descriptor of the file */
    SF_UINT64(file_size)             /* Output. Size of the file */
    SF_UINT64(file_offset)           /* Input/Output. Starting offset */
    SF_INT64(file_bytes)             /* Input/Output. number of bytes to send */
    /* ----- trailer parms ----- */
    void      *trailer_data;         /* Input/Output. Points to trailer buf */
    uint_t     trailer_length;       /* Input/Output. Length of the trailer */
    /* ----- return info ----- */
    SF_UINT64(bytes_sent)            /* Output. number of bytes sent */
};
```

*header\_data*

Points to a buffer that contains header data which is to be sent before the file data. May be a NULL pointer if *header\_length* is 0. This field will be updated by **send\_file** when header is transmitted - that is, *header\_data* + number of bytes of the header sent.

*header\_length*

Specifies the number of bytes in the *header\_data*. This field must be set to 0 to indicate that header data is not to be sent. This field will be updated by **send\_file** when header is transmitted - that is, *header\_length* - number of bytes of the header sent.

*file\_descriptor*

Specifies the file descriptor for a file that has been opened and is readable. This is the descriptor for the file that contains the data to be transmitted. The *file\_descriptor* is ignored when *file\_bytes* = 0. This field is not updated by **send\_file**.

*file\_size*

Contains the byte size of the file specified by *file\_descriptor*. This field is filled in by the kernel.

*file\_offset*

Specifies the byte offset into the file from which to start sending data. This field is updated by the **send\_file** when file data is transmitted - that is, *file\_offset* + number of bytes of the file data sent.

*file\_bytes*

Specifies the number of bytes from the file to be transmitted. Setting *file\_bytes* to -1 transmits the entire file from the *file\_offset*. When this field is not set to -1, it is updated by **send\_file** when file data is transmitted - that is, *file\_bytes* - number of bytes of the file data sent.

*trailer\_data*

Points to a buffer that contains trailer data which is to be sent after the file data. May be a NULL pointer if *trailer\_length* is 0. This field will be updated by **send\_file** when trailer is transmitted - that is, *trailer\_data* + number of bytes of the trailer sent.

*trailer\_length*

Specifies the number of bytes in the *trailer\_data*. This field must be set to 0 to indicate that trailer data is not to be sent. This field will be updated by **send\_file** when trailer is transmitted - that is, *trailer\_length* - number of bytes of the trailer sent.

*bytes\_sent*

Contains number of bytes that were actually sent in this call to **send\_file**. This field is filled in by the kernel.

All fields marked with Input in the *sf\_parms* structure requires setup by an application prior to the **send\_file** calls. All fields marked with Output in the *sf\_parms* structure adjusts by **send\_file** when it successfully transmitted data, that is, either the specified data transmission is partially or completely done.

The **send\_file** subroutine attempts to write *header\_length* bytes from the buffer pointed to by *header\_data*, followed by *file\_bytes* from the file associated with *file\_descriptor*, followed by *trailer\_length* bytes from the buffer pointed to by *trailer\_data*, over the connection associated with the socket pointed to by *Socket\_p*.

As the data is sent, the kernel updates the parameters pointed by *sf\_iobuf* so that if the **send\_file** has to be called multiple times (either due to interruptions by signals, or due to non-blocking I/O mode) in order to complete a file data transmission, the application can reissue the **send\_file** command without setting or re-adjusting the parameters over and over again.

If the application sets *file\_offset* greater than the actual file size, or *file\_bytes* greater than (the actual file size - *file\_offset*), the return value will be -1 with errno EINVAL.

*flags* Specifies the following attributes:

***SF\_CLOSE***

Closes the socket pointed to by *Socket\_p* after the data has been successfully sent or queued for transmission.

***SF\_REUSE***

Prepares the socket for reuse after the data has been successfully sent or queued for transmission and the existing connection closed.

**Note:** This option is currently not supported on this operating system.

***SF\_DONT\_CACHE***

Does not put the specified file in the Network Buffer Cache.

***SF\_SYNC\_CACHE***

Verifies/Updates the Network Buffer Cache for the specified file before transmission.

When the *SF\_CLOSE* flag is set, the connected socket specified by *Socket\_p* will be disconnected and closed by **send\_file** after the requested transmission has been successfully done. The socket descriptor pointed to by *Socket\_p* will be set to -1. This flag won't take effect if **send\_file** returns non-0.

The flag *SF\_REUSE* currently is not supported by AIX. When this flag is specified, the socket pointed by *Socket\_p* will be closed and returned as -1. A new socket needs to be created for the next connection.

**send\_file** will take advantage of a Network Buffer Cache in kernel memory to dynamically cache the output file data. This will help to improve the **send\_file** performance for files which are:

1. accessed repetitively through network and
2. not changed frequently.

Applications can exclude the specified file from being cached by using the *SF\_DONT\_CACHE* flag. **send\_file** will update the cache every so often to make sure that the file data in cache is valid for a certain time period. The network option parameter "send\_file\_duration" controlled by the **no** command can be modified to configure the interval of the **send\_file** cache validation, the default is 300 (in seconds). Applications can use the *SF\_SYNC\_CACHE* flag to ensure that a cache validation of the specified file will occur before the file is sent by **send\_file**, regardless the value of the "send\_file\_duration". Other Network Buffer Cache related parameters are "nbc\_limit", "nbc\_max\_cache", and "nbc\_min\_cache". For additional information, see the **no** command.

## Return Value

There are three possible return values from **send\_file**:

Value	Description
-------	-------------

- |    |   |
|----|---|
| -1 | an error has occurred, <i>errno</i> contains the error code.  |
| 0  | the command has completed successfully.   |
| 1  | the command was completed partially, some data has been transmitted but the command has to return for some reason, for example, the command was interrupted by signals. |

The fields marked with Output in the *sf\_parms* structure (pointed to by *sf\_iobuf*) is updated by **send\_file** when the return value is either 0 or 1. The *bytes\_sent* field contains the total number of bytes that were sent in this call. It is always true that *bytes\_sent* (Output) <= *header\_length*(Input) + *file\_bytes*(Input) + *trailer\_length* (Input).

The **send\_file** supports the blocking I/O mode and the non-blocking I/O mode. In the blocking I/O mode, **send\_file** blocks until all file data (plus the header and the trailer) is sent. It adjusts the *sf\_iobuf* to reflect the transmission results, and return 0. It is possible that **send\_file** can be interrupted before the request is fully done, in that case, it adjusts the *sf\_iobuf* to reflect the transmission progress, and return 1.

In the non-blocking I/O mode, the **send\_file** transmits as much as the socket space allows, adjusts the *sf\_iobuf* to reflect the transmission progress, and returns either 0 or 1. When there is no socket space in

the system to buffer any of the data, the **send\_file** returns -1 and sets **errno** to **EWOULDBLOCK**. **select** or **poll** can be used to determine when it is possible to send more data.

**Possible errno returned:**

<b>EBADF</b>	Either the socket or the file descriptor parameter is not valid.
<b>ENOTSOCK</b>	The socket parameter refers to a file, not a socket.
<b>EPROTONOSUPPORT</b>	Protocol not supported.
<b>EFAULT</b>	The addresses specified in the HeaderTrailer parameter is not in a writable part of the user-address space.
<b>EINTR</b>	The operation was interrupted by a signal before any data was sent. (If some data was sent, <b>send_file</b> returns the number of bytes sent before the signal, and <b>EINTR</b> is not set).
<b>EINVAL</b>	The offset, length of the HeaderTrailer, or flags parameter is invalid.
<b>ENOTCONN</b>	A <b>send_file</b> on a socket that is not connected, a <b>send_file</b> on a socket that has not completed the connect sequence with its peer, or is no longer connected to its peer.
<b>EWOULDBLOCK</b>	The socket is marked non-blocking and the requested operation would block.
<b>ENOMEM</b>	No memory is available in the system to perform the operation.

**PerformanceNote**

By taking advantage of the Network Buffer Cache, **send\_file** provides better performance and network throughput for file transmission. It is recommended for files bigger than 4K bytes.

## Related Information

The **connect** subroutine, **getsockopt** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **sendmsg** subroutine, **sendto** subroutine, **setsockopt** subroutine, **shutdown** subroutine, **socket** subroutine.

Sockets Overview and Understanding Socket Data Transfer in *AIX Version 6.1 Communications Programming Concepts*.

---

## set\_auth\_method Subroutine

### Purpose

Sets the authentication methods for the rcmds for this system.

### Library

Authentication Methods Library (**libauthm.a**)

### Syntax

### Description

This method configures the authentication methods for the system. The authentication methods should be passed to the function in the order in which they should be attempted in the unsigned integer pointer in which the user passed.

The list is an array of unsigned integers terminated by a zero. Each integer identifies an authentication method. The order that a client should attempt to authenticate is defined by the order of the list.

The flags identifying the authentication methods are defined in the **/usr/include/authm.h** file.

Any undefined bits in the input parameter invalidate the entire command. If the same authentication method is specified twice or if any authentication method is specified after Standard AIX, the command fails.

The user must have root authority or this method fails.

## Parameter

*authm* Points to an array of unsigned integers. The list of authentication methods to be set is terminated by a zero.

## Return Values

Upon successful completion, the **set\_auth\_method** subroutine returns a zero.

Upon unsuccessful completion, the **set\_auth\_method** subroutine returns an **errno**.

## Related Information

The **chauthent** command, **ftp** command, **lsauthent** command, **rcp** command, **rlogin** command, **rsh** command, **telnet**, **tn**, or **tn3270** command.

The **get\_auth\_method** subroutine.

Communications and networks in *Networks and communication management*.

Authentication and the secure rcmds in *Networks and communication management*.

---

## setdomainname Subroutine

### Purpose

Sets the name of the current domain.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int setdomainname ( Name, NameLen)
char *Name;
int NameLen;
```

### Description

The **setdomainname** subroutine sets the name of the domain for the host machine. It is normally used when the system is bootstrapped. You must have root user authority to run this subroutine.

The purpose of domains is to enable two distinct networks that may have host names in common to merge. Each network would be distinguished by having a different domain name. At the current time, only Network Information Service (NIS) makes use of domains set by this subroutine.

All applications containing the **setdomainname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

**Note:** Domain names are restricted to 256 characters.

## Parameters

<i>Name</i>	Specifies the domain name to be set.
<i>Namelen</i>	Specifies the size of the array pointed to by the <i>Name</i> parameter.

## Return Values

If the call succeeds, a value of 0 is returned. If the call is unsuccessful, a value of -1 is returned and an error code is placed in the **errno** global variable.

## Error Codes

The following errors may be returned by this subroutine:

Error	Description
<b>EFAULT</b>	The <i>Name</i> parameter gave an invalid address.
<b>EPERM</b>	The caller was not the root user.

## Related Information

The **getdomainname** subroutine, **gethostname** subroutine, **sethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## sethostent Subroutine

### Purpose

Opens network host file.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
**libnis**  
(**liblocal**)

### Syntax

```
#include <netdb.h>
sethostent ( StayOpen)
int StayOpen;
```

### Description

When using the **sethostent** subroutine in DNS/BIND name service resolution, **sethostent** allows a request for the use of a connected socket using TCP for queries. If the *StayOpen* parameter is non-zero, this sets the option to send all queries to the name server using TCP and to retain the connection after each call to **gethostbyname** or **gethostbyaddr**.

When using the **sethostent** subroutine to search the **/etc/hosts** file, **sethostent** opens and rewinds the **/etc/hosts** file. If the *StayOpen* parameter is non-zero, the hosts database is not closed after each call to **gethostbyname** or **gethostbyaddr**.

## Parameters

**StayOpen** When used in NIS name resolution and to search the local **/etc/hosts** file, it contains a value used to indicate whether to close the host file after each call to **gethostbyname** and **gethostbyaddr**. A non-zero value indicates not to close the host file after each call and a zero value allows the file to be closed.

When used in DNS/BIND name resolution, a non-zero value retains the TCP connection after each call to **gethostbyname** and **gethostbyaddr**. A value of zero allows the connection to be closed.

## Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/etc/include/netdb.h</b>	Contains the network database structure.

## Related Information

The **endhostent** subroutine, **gethostbyaddr** subroutine, **gethostbyname** subroutine, **gethostent** subroutine.

Sockets Overview and Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## sethostent\_r Subroutine

### Purpose

Opens network host file.

### Library

Standard C Library (**libc.a**)  
(**libbind**)  
(**libnis**)  
(**liblocal**)

### Syntax

```
#include <netdb.h>
sethostent_r (StayOpenflag, ht_data)
```

```
int StayOpenflag;
struct hostent_data *ht_data;
```

### Description

When using the **sethostent\_r** subroutine in DNS/BIND name service resolution, **sethostent\_r** allows a request for the use of a connected socket using TCP for queries. If the *StayOpen* parameter is non-zero, this sets the option to send all queries to the name server using TCP and to retain the connection after each call to **gethostbyname\_r** or **gethostbyaddr\_r**.

When using the **sethostent\_r** subroutine to search the **/etc/hosts** file, **sethostent\_r** opens and rewinds the **/etc/hosts** file. If the *StayOpen* parameter is non-zero, the hosts database is not closed after each call to **gethostbyname\_r** or **gethostbyaddr\_r**. It internally runs the **sethostent** command.

## Parameters

<i>StayOpenflag</i>	When used in NIS name resolution and to search the local <b>/etc/hosts</b> file, it contains a value used to indicate whether to close the host file after each call to the <b>gethostbyname</b> and <b>gethostbyaddr</b> subroutines. A non-zero value indicates not to close the host file after each call, and a zero value allows the file to be closed.
	When used in DNS/BIND name resolution, a non-zero value retains the TCP connection after each call to <b>gethostbyname</b> and <b>gethostbyaddr</b> . A value of zero allows the connection to be closed.
<i>ht_data</i>	Points to the <b>hostent_data</b> structure.

## Files

<b>/etc/hosts</b>	Contains the host name database.
<b>/etc/netsvc.conf</b>	Contains the name services ordering.
<b>/etc/include/netdb.h</b>	Contains the network database structure.

## Related Information

“endhostent\_r Subroutine” on page 46, “gethostbyname\_r Subroutine” on page 78, “gethostbyaddr\_r Subroutine” on page 75, and “gethostent\_r Subroutine” on page 80.

---

## sethostid Subroutine

### Purpose

Sets the unique identifier of the current host.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int sethostid ( HostID)
int HostID;
```

### Description

The **sethostid** subroutine allows a calling process with a root user ID to set a new 32-bit identifier for the current host. The **sethostid** subroutine enables an application program to reset the host ID.

All applications containing the **sethostid** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*HostID* Specifies the unique 32-bit identifier for the current host.

## Return Values

Upon successful completion, the **sethostid** subroutine returns a value of 0.

If the **sethostid** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

## Error Codes

The **sethostid** subroutine is unsuccessful if the following is true:

Error	Description
<b>EPERM</b>	The calling process did not have an effective user ID of root user.

## Related Information

The **getsockname** subroutine, **gethostid** subroutine, **gethostname** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## sethostname Subroutine

### Purpose

Sets the name of the current host.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int sethostname ( Name, NameLength)
char *Name;
int NameLength;
```

### Description

The **sethostname** subroutine sets the name of a host machine. Only programs with a root user ID can use this subroutine.

The **sethostname** subroutine allows a calling process with root user authority to set the internal host name of a machine on a network.

All applications containing the **sethostname** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Name</i>	Specifies the name of the host machine.
<i>NameLength</i>	Specifies the length of the <i>Name</i> array.

### Return Values

Upon successful completion, the system returns a value of 0.

If the **sethostname** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.

- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

## Error Codes

The **sethostname** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EFAULT</b>	The <i>Name</i> parameter or <i>NameLength</i> parameter gives an address that is not valid.
<b>EPERM</b>	The calling process did not have an effective root user ID.

## Related Information

The **gethostid** subroutine, **gethostname** subroutine, **sethostid** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## setnetent Subroutine

### Purpose

Opens the **/etc/networks** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void setnetent (StayOpen)
int StayOpen;
```

### Description

The **setnetent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **setnetent** subroutine opens the **/etc/networks** file and sets the file marker at the beginning of the file.

All applications containing the **setnetent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>StayOpen</i>	Contains a value used to indicate when to close the <b>/etc/networks</b> file.
	Specifying a value of 0 closes the <b>/etc/networks</b> file after each call to the <b>getnetent</b> subroutine.
	Specifying a nonzero value leaves the <b>/etc/networks</b> file open after each call.

### Return Values

If an error occurs or the end of the file is reached, the **setnetent** subroutine returns a null pointer.

## Files

**/etc/networks**

Contains official network names.

## Related Information

The **endnetent** subroutine, **getnetbyaddr** subroutine, **getnetbyname** subroutine, **getnetent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## setnetent\_r Subroutine

### Purpose

Opens the **/etc/networks** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
int setnetent_r(StayOpenflag, net_data)
struct netent_data *net_data;
int StayOpenflag;
```

### Description

The **setnetent\_r** subroutine opens the **/etc/networks** file and sets the file marker at the beginning of the file.

### Parameters

*StayOpenflag*

Contains a value used to indicate when to close the **/etc/networks** file.

Specifying a value of 0 closes the **/etc/networks** file after each call to the **getnetent** subroutine. Specifying a nonzero value leaves the **/etc/networks** file open after each call.

*net\_data*

Points to the **netent\_data** structure.

## Files

**/etc/networks**

Contains official network names.

## Related Information

“endnetent\_r Subroutine” on page 47, “getnetbyaddr\_r Subroutine” on page 88, “getnetbyname\_r Subroutine” on page 90, and “getnetent\_r Subroutine” on page 92.

---

## setnetgrent\_r Subroutine

### Purpose

Handles the group network entries.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <netdb.h>
int setnetgrent_r(NetGroup,ptr)
char *NetGroup;
void **ptr;
```

## Description

The **setnetgrent\_r** subroutine functions the same as the **setnetgrent** subroutine.

The **setnetgrent\_r** subroutine establishes the network group from which the **getnetgrent\_r** subroutine will obtain members. This subroutine also restarts calls to the **getnetgrent\_r** subroutine from the beginning of the list. If the previous **setnetgrent\_r** call was to a different network group, an **endnetgrent\_r** call is implied. The **endnetgrent\_r** subroutine frees the space allocated during the **getnetgrent\_r** calls.

## Parameters

<i>NetGroup</i>	Points to a network group.
<i>ptr</i>	Keeps the function threadsafe.

## Return Values

The **setnetgrent\_r** subroutine returns a 0 if successful and a -1 if unsuccessful.

## Files

<i>/etc/netgroup</i>	Contains network groups recognized by the system.
<i>/usr/include/netdb.h</i>	Contains the network database structures.

## Related Information

“getnetgrent\_r Subroutine” on page 93, and “endnetgrent\_r Subroutine” on page 48.

---

## setprotoent Subroutine

### Purpose

Opens the **/etc/protocols** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void setprotoent (StayOpen)
int StayOpen;
```

## Description

The **setprotoent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **setprotoent** subroutine opens the **/etc/protocols** file and sets the file marker to the beginning of the file.

All applications containing the **setprotoent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*StayOpen*                      Indicates when to close the **/etc/protocols** file.

                                    Specifying a value of 0 closes the file after each call to **getprotoent**.

                                    Specifying a nonzero value allows the **/etc/protocols** file to remain open after each subroutine.

## Return Values

The return value points to static data that is overwritten by subsequent calls.

## Files

**/etc/protocols**                      Contains the protocol names.

## Related Information

The **endprotoent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## setprotoent\_r Subroutine

### Purpose

Opens the **/etc/protocols** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>

int setprotoent_r(StayOpenflag, proto_data);
int StayOpenflag;
struct protoent_data *proto_data;
```

## Description

The **setprotoent\_r** subroutine opens the **/etc/protocols** file and sets the file marker to the beginning of the file.

## Parameters

*StayOpenflag* Indicates when to close the **/etc/protocols** file.

Specifying a value of 0 closes the file after each call to **getprotoent**. Specifying a nonzero value allows the **/etc/protocols** file to remain open after each subroutine.

## Files

**/etc/protocols** Contains the protocol names.

## Related Information

“endprotoent\_r Subroutine” on page 50, “getprotobyname\_r Subroutine” on page 96, “getprotobyname\_r Subroutine” on page 98, and “getprotoent\_r Subroutine” on page 100.

---

## setservent Subroutine

### Purpose

Opens **/etc/services** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>
void setservent ( StayOpen)
int StayOpen;
```

### Description

The **setservent** subroutine is threadsafe in AIX 4.3 and later. However, the return value points to static data that is overwritten by subsequent calls. This data must be copied to be saved for use by subsequent calls.

The **setservent** subroutine opens the **/etc/services** file and sets the file marker at the beginning of the file.

All applications containing the **setservent** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

## Parameters

*StayOpen* Indicates when to close the **/etc/services** file.

Specifying a value of 0 closes the file after each call to the **getservent** subroutine.

Specifying a nonzero value allows the file to remain open after each call.

## Return Values

If an error occurs or the end of the file is reached, the **setservent** subroutine returns a null pointer.

## Files

**/etc/services**                      Contains service names.

## Related Information

The **endprotoent** subroutine, **endservent** subroutine, **getprotobyname** subroutine, **getprotobynumber** subroutine, **getprotoent** subroutine, **getservbyname** subroutine, **getservbyport** subroutine, **getservent** subroutine, **setprotoent** subroutine.

Sockets Overview and Understanding Network Address Translation in *AIX Version 6.1 Communications Programming Concepts*.

---

## setservent\_r Subroutine

### Purpose

Opens **/etc/services** file and sets the file marker.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <netdb.h>

int setservent_r(StayOpenflag, serv_data)
int StayOpenflag;
struct servent_data serv_data;
```

### Description

The **setservent\_r** subroutine opens the **/etc/services** file and sets the file marker at the beginning of the file.

### Parameters

<i>StayOpenflag</i>	Indicates when to close the <b>/etc/services</b> file.
	Specifying a value of 0 closes the file after each call to the <b>getservent</b> subroutine. Specifying a nonzero value allows the file to remain open after each call.
<i>serv_data</i>	Points to the <b>servent_data</b> structure.

## Files

**/etc/services**                      Contains service names.

## Related Information

“endservent\_r Subroutine” on page 51, “getservbyport\_r Subroutine” on page 106, “getservent\_r Subroutine” on page 109, and “getservbyname\_r Subroutine” on page 104.

---

## setsockopt Subroutine

### Purpose

Sets socket options.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/socketvar.h>
#include <sys/atmsock.h> /*Needed for SOCK_CONN_DGRAM socket type
only*/
```

```
int setsockopt
(Socket, Level, OptionName, OptionValue, OptionLength)
int Socket, Level, OptionName;
const void * OptionValue;
socklen_t OptionLength;
```

### Description

The **setsockopt** subroutine sets options associated with a socket. Options can exist at multiple protocol levels. The options are always present at the uppermost socket level.

The **setsockopt** subroutine provides an application program with the means to control a socket communication. An application program can use the **setsockopt** subroutine to enable debugging at the protocol level, allocate buffer space, control time outs, or permit socket data broadcasts. The **/usr/include/sys/socket.h** file defines all the options available to the **setsockopt** subroutine.

When setting socket options, specify the protocol level at which the option resides and the name of the option.

Use the parameters *OptionValue* and *OptionLength* to access option values for the **setsockopt** subroutine. These parameters identify a buffer in which the value for the requested option or options is returned.

All applications containing the **setsockopt** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the unique socket name.
<i>Level</i>	Specifies the protocol level at which the option resides. To set options at:

#### Socket level

Specifies the *Level* parameter as **SOL\_SOCKET**.

#### Other levels

Supplies the appropriate protocol number for the protocol controlling the option. For example, to indicate that an option will be interpreted by the TCP protocol, set the *Level* parameter to the protocol number of TCP, as defined in the **netinet/in.h** file. Similarly, to indicate that an option will be interpreted by ATM protocol, set the *Level* parameter to **NDDPROTO\_ATM**, as defined in **sys/atmsock.h**.

*OptionName* Specifies the option to set. The *OptionName* parameter and any specified options are passed uninterpreted to the appropriate protocol module for interpretation. The **sys/socket.h** file defines the socket protocol level options. The **netinet/tcp.h** file defines the TCP protocol level options. The socket level options can be enabled or disabled; they operate in a toggle fashion.

The following list defines socket protocol level options found in the **sys/socket.h** file:

#### **SO\_DEBUG**

Turns on recording of debugging information. This option enables or disables debugging in the underlying protocol modules. Set this option in one of the following ways at the command level:

- Use the **sodebug** command, which turns on or off this option for existing sockets.
- Specify `[DEBUG=level]` in the **wait/nowait** field of a service in **inetd.conf** in order to turn on this option for the specific service.
- Set the **sodebug\_env** parameter to `no`, and specify `SODEBUG=level` in the process environment. This turns on or off this option for all subsequent sockets created by the process.

The value for *level* can be either `min`, `normal`, or `detail`.

#### **SO\_REUSEADDR**

Specifies that the rules used in validating addresses supplied by a **bind** subroutine should allow reuse of a local port.

**SO\_REUSEADDR** allows an application to explicitly deny subsequent **bind** subroutine to the port/address of the socket with **SO\_REUSEADDR** set. This allows an application to block other applications from binding with the **bind** subroutine.

#### **SO\_REUSEPORT**

Specifies that the rules used in validating addresses supplied by a **bind** subroutine should allow reuse of a local port/address combination. Each binding of the port/address combination must specify the **SO\_REUSEPORT** socket option

#### **SO\_CKSUMREV**

Enables performance enhancements in the protocol layers. If the protocol supports this option, enabling causes the protocol to defer checksum verification until the user's data is moved into the user's buffer (on **recv**, **recvfrom**, **read**, or **recvmsg** thread). This can cause applications to be awakened when no data is available, in the case of a checksum error. In this case, **EAGAIN** is returned. Applications that set this option must handle the **EAGAIN** error code returned from a receive call.

#### **SO\_KEEPAIVE**

Monitors the activity of a connection by enabling or disabling the periodic transmission of ACK messages on a connected socket. The idle interval time can be designated using the TCP/IP **no** command. Broken connections are discussed in "Understanding Socket Types and Protocols" in *AIX Version 6.1 Communications Programming Concepts*.

#### **SO\_DONTRROUTE**

Does not apply routing on outgoing messages. Indicates that outgoing messages should bypass the standard routing facilities. Instead, they are directed to the appropriate network interface according to the network portion of the destination address.

**SO\_BROADCAST**

Permits sending of broadcast messages.

**SO\_LINGER**

Lingers on a **close** subroutine if data is present. This option controls the action taken when an unsent messages queue exists for a socket, and a process performs a **close** subroutine on the socket.

If **SO\_LINGER** is set, the system blocks the process during the **close** subroutine until it can transmit the data or until the time expires. If **SO\_LINGER** is not specified and a **close** subroutine is issued, the system handles the call in a way that allows the process to continue as quickly as possible.

The **sys/socket.h** file defines the **linger** structure that contains the **l\_linger** member for specifying linger time interval. If linger time is set to anything but 0, the system tries to send any messages queued on the socket. The maximum value that the **l\_linger** member can be set to is 65535. If the application has requested SPEC1170 compliant behavior by exporting the **XPG\_SUS\_ENV** environment variable, the linger time is *n* seconds; otherwise, the linger time is *n*/100 seconds (ticks), where *n* is the value of the **l\_linger** member.

**SO\_OOINLINE**

Leaves received out-of-band data (data marked urgent) in line.

**SO\_SNDBUF**

Sets send buffer size.

**SO\_RCVBUF**

Sets receive buffer size.

**SO\_SNDLOWAT**

Sets send low-water mark.

**SO\_RCVLOWAT**

Sets receive low-water mark.

**SO\_SNDTIMEO**

Sets send time out. This option is settable, but currently not used.

**SO\_RCVTIMEO**

Sets receive time out. This option is settable, but currently not used.

**SO\_ERROR**

Sets the retrieval of error status and clear.

**SO\_TYPE**

Sets the retrieval of a socket type.

The following list defines TCP protocol level options found in the **netinet/tcp.h** file:

**TCP\_KEEPCNT**

Specifies the maximum number of keepalive packets to be sent to validate a connection. This socket option value is inherited from the parent socket. The default is 8.

**TCP\_KEEPIDLE**

Specifies the number of seconds of idle time on a connection after which TCP sends a keepalive packet. This socket option value is inherited from the parent socket from the accept system call. The default value is 7200 seconds (14400 half seconds).

**TCP\_KEEPINTVL**

Specifies the interval of time between keepalive packets. It is measured in seconds. This socket option value is inherited from the parent socket from the accept system call. The default value is 75 seconds (150 half seconds).

**TCP\_NODELAY**

Specifies whether TCP should follow the Nagle algorithm for deciding when to send data. By default, TCP will follow the Nagle algorithm. To disable this behavior, applications can enable **TCP\_NODELAY** to force TCP to always send data immediately. For example, **TCP\_NODELAY** should be used when there is an application using TCP for a request/response.

**TCP\_RFC1323**

Enables or disables RFC 1323 enhancements on the specified TCP socket. An application might contain the following lines to enable RFC 1323:

```
int on=1;
setsockopt(s, IPPROTO_TCP, TCP_RFC1323, &on, sizeof(on));
```

**TCP\_STDURG**

Enables or disables RFC 1122 compliant urgent point handling. By default, TCP implements urgent pointer behavior compliant with the 4.2 BSD operating system, i.e., this option defaults to 0.

**TCP\_NODELAYACK**

Specifies if TCP needs to send immediate acknowledgement packets to the sender. If this option is not set, TCP delays sending the acknowledgement packets by up to 200 ms. This allows the acknowledgements to be sent along with the data on a response and minimizes system overhead. Setting this TCP option might cause a slight increase in system overhead, but can result in higher performance for network transfers if the sender is waiting on the receiver's acknowledgements.

Beginning at AIX 4.3.2, TCP protocol level socket options are inherited from listening sockets to new sockets. Prior to 4.3.2, only the TCP\_RFC1323 option was inherited.

The following list defines ATM protocol level options found in the **sys/atmsock.h** file:

**SO\_ATM\_PARAM**

Sets all ATM parameters. This socket option can be used instead of using individual sockets options described below. It uses the **connect\_ie** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_AAL\_PARM**

Sets ATM AAL(Adaptation Layer) parameters. It uses the **aal\_parm** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_TRAFFIC\_DES**

Sets ATM Traffic Descriptor values. It uses the **traffic** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BEARER**

Sets ATM Bearer capability. It uses the **bearer** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BHLI**

Sets ATM Broadband High Layer Information. It uses the **bhli** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_BLLI**

Sets ATM Broadband Low Layer Information. It uses the **blli** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_QOS**

Sets ATM Quality Of Service values. It uses the **qos\_parm** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_TRANSIT\_SEL**

Sets ATM Transit Selector Carrier. It uses the **transit\_sel** structure defined in **sys/call\_ie.h** file.

**SO\_ATM\_ACCEPT**

Indicates acceptance of an incoming ATM call, which was indicated to the application via *ACCEPT* system call. This must be issues for the incoming connection to be fully established. This allows negotiation of ATM parameters.

**SO\_ATM\_MAX\_PEND**

Sets the number of outstanding transmit buffers that are permitted before an error indication is returned to applications as a result of a transmit operation. This option is only valid for non best effort types of virtual circuits. *OptionValue/OptionLength* point to a byte which contains the value that this parameter will be set to.

*OptionValue* The *OptionValue* parameter takes an *Int* parameter. To enable a Boolean option, set the *OptionValue* parameter to a nonzero value. To disable an option, set the *OptionValue* parameter to 0.

The following options enable and disable in the same manner:

- **SO\_DEBUG**
- **SO\_REUSEADDR**
- **SO\_KEEPAIVE**
- **SO\_DONTROUTE**
- **SO\_BROADCAST**
- **SO\_OOINLINE**
- **SO\_LINGER**
- **TCP\_RFC1323**

*OptionLength* The *OptionLength* parameter contains the size of the buffer pointed to by the *OptionValue* parameter.

Options at other protocol levels vary in format and name.

IP level (**IPPROTO\_IP** level) options are defined as follows:

<b>IP_DONTFRAG</b>	Sets DF bit from now on for every packet in the IP header. Beginning with AIX 5.3, to detect decreases in Path MTU, UDP applications will always need to set this option.
<b>IP_FINDPMTU</b>	Sets enable/disable PMTU discovery for this path. Protocol level path MTU discovery should be enabled for the discovery to happen.
<b>IP_PMTUAGE</b>	Sets the age of PMTU. Specifies the frequency of PMT reductions discovery for the session. Setting it to 0 (zero) implies infinite age and PMTU reduction discovery will not be attempted. This will replace the previously set PMTU age. The new PMTU age will become effective after the currently set timer expires. Beginning with AIX 5.3, this option is unused because UDP applications will always need to set the <b>IP_DONTFRAG</b> socket option to detect decreases in PMTU immediately.
<b>IP_TTL</b>	Sets the time-to-live field in the IP header for every packet. However, for raw sockets, the default MAXTTL value will be used while sending the messages irrespective of the value set using the <b>setsockopt</b> subroutine.
<b>IP_HDRINCL</b>	This option allows users to build their own IP header. It indicates that the complete IP header is included with the data and can be used only for raw sockets.
<b>IP_ADD_MEMBERSHIP</b>	Joins a multicast group as specified in the <i>OptionValue</i> parameter of the <b>ip_mreq</b> structure type.
<b>IP_DROP_MEMBERSHIP</b>	Leaves a multicast group as specified in the <i>OptionValue</i> parameter of the <b>ip_mreq</b> structure type.
<b>IP_MULTICAST_IF</b>	Permits sending of multicast messages on an interface as specified in the <i>OptionValue</i> parameter of the <b>ip_addr</b> structure type. An address of <b>INADDR_ANY</b> (0x00000000) removes the previous selection of an interface in the multicast options. If no interface is specified, the interface leading to the default route is used.
<b>IP_MULTICAST_LOOP</b>	Sets multicast loopback, determining whether or not transmitted messages are delivered to the sending host. An <i>OptionValue</i> parameter of the <b>char</b> type controls the loopback to be on or off.
<b>IP_MULTICAST_TTL</b>	Sets the time-to-live (TTL) for multicast packets. An <i>OptionValue</i> parameter of the <b>char</b> type sets the value of TTL ranging from 0 through 255.
<b>IP_BLOCK_SOURCE</b>	Blocks data from a given source to a given group.
<b>IP_UNBLOCK_SOURCE</b>	Unblocks a blocked source (to undo the <b>IP_BLOCK_SOURCE</b> operation).
<b>IP_ADD_SOURCE_MEMBERSHIP</b>	Joins a source-specific multicast group. If the host is a member of the group, accept data from the source; otherwise, join the group and accept data from the given source.
<b>IP_DROP_SOURCE_MEMBERSHIP</b>	Leaves a source-specific multicast group. Drops the source from the given multicast group list. To drop all sources of a given group, use the <b>IP_DROP_MEMBERSHIP</b> socket option.

IPV6 level (**IPPROTO\_IPV6** level) options are defined as follows:

<b>IPV6_V6ONLY</b>	Restricts AF_INET6 sockets to IPV6 communications only.	
	<b>Option Type:</b>	int (boolean interpretation)

IPV6_UNICAST_HOPS	Allows the user to set the outgoing hop limit for unicast IPV6 packets.	
	<b>Option Type:</b>	int (x)
	<b>Option Value:</b>	x < -1      Error EINVAL
		x == -1      Use kernel default
		0 <= x <= 255      Use x
		x >= 256      Error EINVAL
IPV6_MULTICAST_HOPS	Allows the user to set the outgoing hop limit for multicast IPV6 packets.	
	<b>Option Type:</b>	int (x)
	<b>Option Value:</b>	Interpretation is the same as <b>IPV6_UNICAST_HOPS</b> (listed above).
IPV6_MULTICAST_IF	Allows the user to specify the interface being used for outgoing multicast packets. If specified as 0, the system selects the outgoing interface.	
	<b>Option Type:</b>	unsigned int (index of interface to use)
IPV6_MULTICAST_LOOP	If a multicast datagram is sent to a group that the sending host belongs to, a copy of the datagram is looped back by the IP layer for local delivery (if the option is set to 1). If the option is set to 0, a copy is not looped back.	
	<b>Option Type:</b>	unsigned int
IPV6_JOIN_GROUP	Joins a multicast group on a specified local interface. If the interface index is specified as 0, the kernel chooses the local interface.	
	<b>Option Type:</b>	struct <b>ipv6_mreq</b> as defined in the <b>netinet/in.h</b> file
IPV6_LEAVE_GROUP	Leaves a multicast group on a specified interface.	
	<b>Option Type:</b>	struct <b>ipv6_mreq</b> as defined in the <b>netinet/in.h</b> file
IPV6_CHECKSUM	Specifies that the kernel computes checksums over the data and the pseudo-IPv6 header for a raw socket. The kernel will compute the checksums for outgoing packets as well as verify checksums for incoming packets on that socket. Incoming packets with incorrect checksums will be discarded. This option is disabled by default.	
	<b>Option Type:</b>	int
	<b>Option Value:</b>	Offsets into the user data where the checksum result must be stored. This must be a positive even value. Setting the value to -1 will disable the option.
IPV6_RECVPKTINFO	Causes the destination IPv6 address and arriving interface index of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
IPV6_RECVHOPLIMIT	Causes the hop limit of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)
IPV6_RECVTCLASS	Causes the traffic class of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	<b>Option Type:</b>	int (boolean interpretation)

IPV6_RECVRTHDR	Causes the routing header (if any) of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	Option Type:	int (boolean interpretation)
IPV6_RECVHOPOPTS	Causes the hop-by-hop options header (if any) of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	Option Type:	int (boolean interpretation)
IPV6_RECVDSTOPTS	Causes the destination options header (if any) of incoming IPv6 packets to be received as ancillary data on UDP and raw sockets.	
	Option Type:	int (boolean interpretation)
IPV6_PKTINFO	Sets the source IPv6 address and outgoing interface index for all IPv6 packets being sent on this socket. This option can be cleared by doing a regular <b>setsockopt</b> with <b>ipi6_addr</b> being <b>in6addr_any</b> and <b>ipi6_ifindex</b> being 0.	
	Option Type:	struct in6_pktinfo defined in the <b>netinet/in.h</b> file.
IPV6_NEXTHOP	Sets the next hop for outgoing IPv6 datagrams on this socket. This option can be cleared by doing a regular <b>setsockopt</b> with a 0 length. Note that a memory pointer must still be supplied for the option value in this case.	
	Option Type:	struct sockaddr_in6 defined in the <b>netinet/in.h</b> file.
IPV6_TCLASS	Sets the traffic class for outgoing IPv6 datagrams on this socket. To clear this option, the application can specify -1 as the value.	
	Option Type:	int (x)
	Option Value:	x < -1 Error EINVAL x == -1 Use kernel default 0 <= x <= 255 Use x x >= 256 Error EINVAL
IPV6_RTHDR	Sets the routing header to be used for outgoing IPv6 datagrams on this socket. This option can be cleared by doing a regular <b>setsockopt</b> with a 0 length. Note that a memory pointer must still be supplied for the option value in this case.	
	Option Type:	struct ip6_rthdr defined in the <b>netinet/ip6.h</b> file.
IPV6_HOPOPTS	Sets the hop-by-hop options header to be used for outgoing IPv6 datagrams on this socket. This option can be cleared by doing a regular <b>setsockopt</b> with a 0 length. Note that a memory pointer must still be supplied for the option value in this case.	
	Option Type:	struct ip6_hbh defined in the <b>netinet/ip6.h</b> file.
IPV6_DSTOPTS	Sets the destination options header to be used for outgoing IPv6 datagrams on this socket. This header will follow a routing header (if present) and will also be used when there is no routing header specified. This option can be cleared by doing a regular <b>setsockopt</b> with a 0 length. Note that a memory pointer must still be supplied for the option value in this case.	
	Option Type:	struct ip6_dest defined in the <b>netinet/ip6.h</b> file.
IPV6_RTHDRDSTOPTS	Sets the destination options header to be used for outgoing IPv6 datagrams on this socket. This header will precede a routing header (if present). If no routing header is specified, this option will be silently ignored. This option can be cleared by doing a regular <b>setsockopt</b> with a 0 length. Note that a memory pointer must still be supplied for the option value in this case.	
	Option Type:	struct ip6_dest defined in the <b>netinet/ip6.h</b> file.

IPV6_USE_MIN_MTU	Sets this option to control IPv6 path MTU discovery.	
	Option Type:	int
	Option Value:	-1 Performs path MTU discovery for unicast destinations, but does not perform it for multicast destinations. 0 Always performs path MTU discovery. 1 Always disables path MTU discovery and sends packets at the minimum MTU.
IPV6_DONTFRAG	Setting this option prevents fragmentation of outgoing IPv6 packets on this socket. If a packet is being sent that is larger than the outgoing interface MTU, the packet will be discarded.	
	Option Type:	int (boolean interpretation)
IPV6_RECVPATHMTU	Enables the receipt of IPV6_PATHMTU ancillary data items by setting this option.	
	Option Type:	int (boolean interpretation)
MCAST_JOIN_GROUP	Joins the multicast group as specified in the <i>OptionValue</i> parameter of the <b>group_req</b> structure. If the specified interface index is 0, the kernel chooses the default interface.	
	Option Type:	struct group_req as defined in the <b>netinet/in.h</b> file
MCAST_LEAVE_GROUP	Leaves the multicast group as specified in the <i>OptionValue</i> parameter of the <b>group_req</b> structure.	
	Option Type:	struct group_req as defined in the <b>netinet/in.h</b> file
MCAST_BLOCK_SOURCE	Blocks data from the specified source to the specified multicast group.	
	Option Type:	struct group_source_req as defined in the <b>netinet/in.h</b> file
MCAST_UNBLOCK_SOURCE	Unblocks data from the specified source to the specified multicast group. The option is used to undo the <b>MCAST_BLOCK_SOURCE</b> operation.	
	Option Type:	struct group_source_req as defined in the <b>netinet/in.h</b> file
MCAST_JOIN_SOURCE_GROUP	Joins a source-specific multicast group. If the host is already a member of the group, accept data from the specified source; otherwise, join the group and accept data from the specified source.	
	Option Type:	struct group_source_req as defined in the <b>netinet/in.h</b> file
MCAST_LEAVE_SOURCE_GROUP	Leaves a source-specific multicast group. Leaves the specified source from the specified multicast group. To leave all sources of the multicast group, use the <b>IPV6_LEAVE_GROUP</b> or <b>MCAST_LEAVE_GROUP</b> socket option.	
	Option Type:	struct group_source_req as defined in the <b>netinet/in.h</b> file

ICMPV6 level (**IPPROTO\_ICMPV6** level) options are defined as follows:

ICMP6_FILTER	Allows the user to filter ICMPV6 messages by the ICMPV6 type field. In order to clear an existing filter, issue a <b>setsockopt</b> call with zero length.	
	Option Type:	The <b>icmp6_filter</b> structure defined in the <b>netinet/icmp6.h</b> file.

The following values (defined in the **/usr/include/netint/tcp.h** file) are used by the **setsockopt** subroutine to configure the **dacinet** functions.

**Note:** The DACinet facility is available only in a CAPP/EAL4+ configured AIX system.

```
tcp.h:#define TCP_ACLFLUSH    0x21    /* clear all DACinet ACLs */
tcp.h:#define TCP_ACLCLEAR   0x22    /* clear DACinet ACL */
tcp.h:#define TCP_ACLADD     0x23    /* Add to DACinet ACL */
tcp.h:#define TCP_ACLDEL     0x24    /* Delete from DACinet ACL */
tcp.h:#define TCP_ACLLS      0x25    /* List DACinet ACL */
tcp.h:#define TCP_ACLBIND    0x26    /* Set port number for TCP_ACLLS */
tcp.h:#define TCP_ACLGID     0x01    /* id being added to ACL is a gid */
tcp.h:#define TCP_ACLUID     0x02    /* id being added to ACL is a gid */
tcp.h:#define TCP_ACLSUBNET   0x04    /* address being added to ACL is a subnet */
tcp.h:#define TCP_ACLDENY    0x08    /* this ACL entry is for denying access */
```

## Return Values

Upon successful completion, a value of 0 is returned.

If the **setsockopt** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

## Error Codes

The **setsockopt** subroutine is unsuccessful if any of the following errors occurs:

<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINVAL</b>	The <i>OptionValue</i> parameter or the <i>OptionLength</i> parameter is invalid or the socket has been shutdown.
<b>ENOBUFS</b>	There is insufficient memory for an internal data structure.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>ENOPROTOOPT</b>	The option is unknown.
<b>EOPNOTSUPP</b>	The option is not supported by the socket family or socket type.

## Examples

To mark a socket for broadcasting:

```
int on=1;
setsockopt(s, SOL_SOCKET, SO_BROADCAST, &on, sizeof(on));
```

To turn on the **TCP\_NODELAYACK** option, run the following:

```
int on=1;
setsockopt(s, IPPROTO_TCP, TCP_NODELAYACK, &on, sizeof(on));
```

## Related Information

The **no** command.

The **bind** subroutine, **endprotoent** subroutine, **getprotobyname** subroutine, **getprotoent** subroutine, **getsockopt** subroutine, **setprotoent** subroutine, **socket** subroutine.

Sockets Overview, Understanding Socket Options, Understanding Socket Types and Protocols in *AIX Version 6.1 Communications Programming Concepts*.

---

## shutdown Subroutine

### Purpose

Shuts down all socket send and receive operations.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int shutdown ( Socket, How)
int Socket, How;
```

### Description

The **shutdown** subroutine disables all receive and send operations on the specified socket.

All applications containing the **shutdown** subroutine must be compiled with the **\_BSD** macro set to a specific value. Acceptable values are 43 and 44. In addition, all socket applications must include the BSD **libbsd.a** library.

### Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>How</i>	Specifies the type of subroutine shutdown. Use the following values:
0	Disables further receive operations.
1	Disables further send operations.
2	Disables further send operations and receive operations.

### Return Values

Upon successful completion, a value of 0 is returned.

If the **shutdown** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

### Error Codes

The **shutdown** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EINVAL</b>	The <i>How</i> parameter is invalid.
<b>ENOTCONN</b>	The socket is not connected.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.

## Files

`/usr/include/sys/socket.h`

Contains socket definitions.

`/usr/include/sys/types.h`

Contains definitions of unsigned data types.

## Related Information

The **getsockopt** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **read** subroutine, **select** subroutine, **send** subroutine, **sendto** subroutine, **setsockopt** subroutine, **socket** subroutine, **write** subroutine.

Sockets Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## SLPAttrCallback Subroutine

### Purpose

Returns the same callback type as the **SLPFindAttrs()** function.

### Syntax

```
typedef SLPBoolean SLPAttrCallback(SLPHandle hSLP,  
                                   const char* pcAttrList,  
                                   SLPErrCode errCode,  
                                   void *pvCookie);
```

### Description

The **SLPAttrCallback** type is the type of the callback function parameter to the **SLPFindAttrs()** function.

The *pcAttrList* parameter contains the requested attributes as a comma-separated list (or is empty if no attributes matched the original tag list).

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> used to initiate the operation.
<i>pcAttrList</i>	A character buffer containing a comma-separated, null-terminated list of attribute ID/value assignments, in SLP wire format: "(attr-id=attr-value-list)"
<i>errCode</i>	An error code indicating if an error occurred during the operation. The callback should check this error code before processing the parameters. If the error code is other than <b>SLP_OK</b> , then the API library can choose to terminate the outstanding operation.
<i>pvCookie</i>	Memory passed down from the client code that called the original API function, starting the operation. Can be NULL.

### Return Values

The client code should return **SLP\_TRUE** if more data is desired; otherwise **SLP\_FALSE** is returned.

## Related Information

"SLPClose Subroutine" on page 228, "SLPEscape Subroutine" on page 229, "SLPFindAttrs Subroutine" on page 230, "SLPFindScopes Subroutine" on page 231, "SLPFindSrvs Subroutine" on page 232, "SLPFindSrvTypes Subroutine" on page 233, "SLPFree Subroutine" on page 234, "SLPGetProperty Subroutine" on page 234, "SLPOpen Subroutine" on page 235, "SLPParseSrvURL Subroutine" on page 236, "SLPSrvTypeCallback Subroutine" on page 239, "SLPSrvURLCallback Subroutine" on page 240, "SLPUnescape Subroutine" on page 240

The Service Location Protocol (SLP) API.

Frees all resources associated with the handle.

```
void SLPClose(SLPHandle hSLP):
```

The **SLPClose** subroutine frees all resources associated with the handle. If the handle was invalid, the function returns silently. Any outstanding synchronous or asynchronous operations are cancelled so that their callback functions will not be called any further.

*hSLP* The **SLPHandle** handle returned from a call to **SLPOpen()**.

“SLPAttrCallback Subroutine” on page 227, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

## The /etc/slp.conf File.

The Service Location Protocol (SLP) API.

Deregisters the advertisement for URL in all scopes and locales.

```
SLPError SLPReg(hSLP, pcURL, callback, pvCookie)
SLPHandle hSLP;
const char *pcURL;
SLPRegReport callback;
void *pvCookie;
```

The **SLPDereg** subroutine deregisters the advertisement for the URL specified by the *pcURL* parameter in all scopes where the service is registered and in all language locales. The deregistration is not confined to the **SLPHandle** locale. Deregistration takes place in all locales.

## Parameters

<i>hSLP</i>	The language-specific <b>SLPHandle</b> handle used for deregistration of services.
<i>pcURL</i>	The URL that needs to be deregistered.
<i>callback</i>	A callback function through which the results of the operation are reported.
<i>pvCookie</i>	The memory passed to callback code from the client. The parameter can be set to NULL.

## Return Values

<b>SLP_OK</b>	The subroutine has run successfully.
<b>SLPError</b>	An error occurred.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine,” “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240, “SLPReg Subroutine” on page 237, “SLPRegReport Callback Subroutine” on page 238.

The **/etc/slp.conf** File.

The Service Location Protocol (SLP) APIs.

---

## SLPEscape Subroutine

### Purpose

Processes an input string and escapes any characters reserved for SLP.

### Syntax

```
SLPError SLPEscape(const char* pcInbuf,  
                  char** ppcOutBuf,  
                  SLPBoolean isTag);
```

### Description

The **SLPEscape** subroutine processes the input string in *pcInbuf* and escapes any characters reserved for SLP. If the *isTag* parameter is **SLPTrue**, **SLPEscape** looks for bad tag characters and signals an error if any are found by returning the **SLP\_PARSE\_ERROR** code. The results are put into a buffer allocated by the API library and returned in the *ppcOutBuf* parameter. This buffer should be deallocated using **SLPFree()** when the memory is no longer needed.

## Parameters

<i>pcInbuf</i>	Pointer to the input buffer to process for escape characters.
<i>ppcOutBuf</i>	Pointer to a pointer for the output buffer with the characters reserved for SLP escaped. Must be freed using <b>SLPFree()</b> when the memory is no longer needed.
<i>isTag</i>	When true, the input buffer is checked for bad tag characters.

## Return Values

The **SLPEscape** subroutine returns **SLP\_PARSE\_ERROR** if any characters are bad tag characters and the *isTag* flag is true; otherwise, it returns **SLP\_OK**, or the appropriate error code if another error occurs.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPFindAttrs Subroutine,” “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The */etc/slp.conf* File.

The Service Location Protocol (SLP) API.

---

## SLPFindAttrs Subroutine

### Purpose

Returns service attributes that match the attribute IDs for the indicated service URL or service type.

### Syntax

```
SLPError SLPFindAttrs(SLPHandle hSLP,
                      const char *pcURLOrServiceType,
                      const char *pcScopeList,
                      const char *pcAttrIds,
                      SLPAttrCallback callback,
                      void *pvCookie);
```

### Description

The **SLPFindAttrs** subroutine returns service attributes matching the attribute IDs for the indicated service URL or service type. If **pcURLOrServiceType** is a service URL, the attribute information returned is for that particular advertisement in the language locale of the *SLPHandle*.

If **pcURLOrServiceType** is a service type name (including naming authority if any), then the attributes for all advertisements of that service type are returned regardless of the language of registration. Results are returned through the *callback*.

The result is filtered with an SLP attribute request filter string parameter. If the filter string is the empty string (""), all attributes are returned.

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> on which to search for attributes.
<i>pcURLOrServiceType</i>	The service URL or service type. Cannot be the empty string.
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. Cannot be the empty string, "".
<i>pcAttrIds</i>	The filter string indicating which attribute values to return. Use the empty string ("" to indicate all values. Wildcards matching all attribute IDs having a particular prefix or suffix are also possible.
<i>callback</i>	A callback function through which the results of the operation are reported.
<i>pvCookie</i>	Memory passed to the callback code from the client. Can be NULL.

## Return Values

If **SLPFindAttrs** is successful, it returns **SLP\_OK**. If an error occurs in starting the operation, one of the **SLPError** codes is returned.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindScopes Subroutine,” “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## SLPFindScopes Subroutine

### Purpose

Sets the *ppcScopeList* parameter to point to a comma-separated list that includes all available scope values.

### Syntax

```
SLPError SLPFindScopes(SLPHandle hSLP,  
                      char** ppcScopeList);
```

### Description

The **SLPFindScopes** subroutine sets the *ppcScopeList* parameter to point to a comma-separated list that includes all available scope values. If there is any order to the scopes, preferred scopes are listed before less desirable scopes. There is always at least one name in the list, the default scope, **DEFAULT**.

### Parameters

*hSLP*

The **SLPHandle** on which to search for scopes.

*ppcScopeList*

A pointer to a char pointer into which the buffer pointer is placed upon return. The buffer is null terminated. The memory should be freed by calling **SLPFree()**.

### Return Values

If no error occurs, **SLPFindScopes** returns **SLP\_OK**; otherwise, it returns the appropriate error code.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## SLPFindSrvs Subroutine

### Purpose

Issues the query for services on the language-specific **SLPHandle** and returns the results through the *callback*.

### Syntax

```
SLPError SLPFindSrvs(SLPHandle hSLP,  
                    const char *pcServiceType,  
                    const char *pcScopeList,  
                    const char *pcSearchFilter,  
                    SLPsrvURLCallback callback,  
                    void *pvCookie);
```

### Description

The **SLPFindSrvs** subroutine issues the query for services on the language-specific **SLPHandle** and returns the results through the callback. The parameters determine the results

### Parameters

<i>hSLP</i>	The language-specific <b>SLPHandle</b> on which to search for services.
<i>pcServiceType</i>	The Service Type String, including authority string if any, for the request, which can be discovered using <b>SLPSrvTypes()</b> . This could be, for example, "service:printer:lpr" or "service:nfs". This cannot be the empty string ("").
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names. This cannot be the empty string ("").
<i>pcSearchFilter</i>	A query formulated of attribute pattern matching expressions in the form of a LDAPv3 Search Filter. If this filter is empty (""), all services of the requested type in the specified scopes are returned.
<i>callback</i>	A callback function through which the results of the operation are reported.
<i>pvCookie</i>	Memory passed to the callback code from the client. Can be NULL.

### Return Values

If **SLPFindSrvs** is successful, it returns **SLP\_OK**. If an error occurs in starting the operation, one of the **SLPError** codes is returned.

### Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The /etc/slp.conf File.

The Service Location Protocol (SLP) API.

---

## SLPFindSrvTypes Subroutine

### Purpose

Issues an SLP service type request.

### Syntax

```
SLPError SLPFindSrvTypes(SLPHandle    hSLP,
                        const char *pcNamingAuthority,
                        const char *pcScopeList,
                        SLPsrvTypeCallback callback,
                        void *pvCookie);
```

### Description

The **SLPFindSrvType()** subroutine issues an SLP service type request for service types in the scopes indicated by the **pcScopeList**. The results are returned through the *callback* parameter. The service types are independent of language locale, but only for services registered in one of the scopes and for the naming authority indicated by *pcNamingAuthority*.

If the naming authority is "\*", then results are returned for all naming authorities. If the naming authority is the empty string, "", then the default naming authority, "IANA", is used. "IANA" is not a valid naming authority name, and it returns a **PARAMETER\_BAD** error when it is included explicitly.

The service type names are returned with the naming authority intact. If the naming authority is the default (that is, the empty string), then it is omitted, as is the separating ".". Service type names from URLs of the **service:** scheme are returned with the "service:" prefix intact.

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> on which to search for types.
<i>pcNamingAuthority</i>	The naming authority to search. Use "*" for all naming authorities and the empty string, "", for the default naming authority.
<i>pcScopeList</i>	A pointer to a char containing a comma-separated list of scope names to search for service types. Cannot be the empty string, "".
<i>callback</i>	A callback function through which the results of the operation are reported.
<i>pvCookie</i>	Memory passed to the callback code from the client. Can be NULL.

### Return Values

If **SLPFindSrvTypes** is successful, it returns **SLP\_OK**. If an error occurs in starting the operation, one of the **SLPError** codes is returned.

### Related Information

"SLPAttrCallback Subroutine" on page 227, "SLPClose Subroutine" on page 228, "SLPEscape Subroutine" on page 229, "SLPFindAttrs Subroutine" on page 230, "SLPFindScopes Subroutine" on page 231, "SLPFindSrvs Subroutine" on page 232, "SLPFree Subroutine" on page 234, "SLPGetProperty Subroutine" on page 234, "SLPOpen Subroutine" on page 235, "SLPParseSrvURL Subroutine" on page 236, "SLPSrvTypeCallback Subroutine" on page 239, "SLPSrvURLCallback Subroutine" on page 240, "SLPUnescape Subroutine" on page 240

The /etc/slp.conf File.

The Service Location Protocol (SLP) API.

---

## SLPFree Subroutine

### Purpose

Frees memory returned from **SLPParseSrvURL()**, **SLPFindScopes()**, **SLPEscape()**, and **SLPUnescape()**.

### Syntax

```
void SLPFree(void* pvMem);
```

### Description

The **SLPFree** subroutine frees memory returned from **SLPParseSrvURL()**, **SLPFindScopes()**, **SLPEscape()**, and **SLPUnescape()**.

### Parameters

<i>pvMem</i>	A pointer to the storage allocated by the <b>SLPParseSrvURL()</b> , <b>SLPEscape()</b> , <b>SLPUnescape()</b> , or <b>SLPFindScopes()</b> function. Ignored if NULL.
--------------	--

### Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPGetProperty Subroutine,” “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## SLPGetProperty Subroutine

### Purpose

Returns the value of the corresponding SLP property name.

### Syntax

```
const char* SLPGetProperty(const char* pcName);
```

### Description

The **SLPGetProperty** subroutine returns the value of the corresponding SLP property name. The returned string is owned by the library and *must not* be freed.

### Parameters

<i>pcName</i>	Null-terminated string with the property name.
---------------	--

## Return Values

If no error, the **SLPGetProperty** subroutine returns a pointer to a character buffer containing the property value. If the property was not set, the subroutine returns the default value. If an error occurs, it returns NULL. The returned string *must not* be freed.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPOpen Subroutine,” “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The /etc/slp.conf File.

The Service Location Protocol (SLP) API.

---

## SLPOpen Subroutine

### Purpose

Returns an **SLPHandle** handle that encapsulates the language locale for SLP requests.

### Syntax

```
SLPError SLPOpen(const char *pcLang, SLPBoolean isAsync, SLPHandle
    *phSLP);
```

### Description

The **SLPOpen** subroutine returns an **SLPHandle** handle in the *phSLP* parameter for the language locale passed in as the *pcLang* parameter. The client indicates if operations on the handle are to be synchronous or asynchronous through the *isAsync* parameter. The handle encapsulates the language locale for SLP requests issued through the handle, and any other resources required by the implementation. However, SLP properties are not encapsulated by the handle; they are global. The return value of the function is an **SLPError** code indicating the status of the operation. Upon failure, the *phSLP* parameter is NULL.

### Implementation Specifics

An **SLPHandle** can only be used for one SLP API operation at a time. If the original operation was started asynchronously, any attempt to start an additional operation on the handle while the original operation is pending results in the return of an **SLP\_HANDLE\_IN\_USE** error from the API function. The **SLPClose()** API function terminates any outstanding calls on the handle. If an implementation is unable to support an asynchronous (resp. synchronous) operation, because of memory constraints or lack of threading support, the **SLP\_NOT\_IMPLEMENTED** flag might be returned when the *isAsync* flag is **SLP\_TRUE** (resp. **SLP\_FALSE**).

### Parameters

*pcLang*

A pointer to an array of characters (AIX supports “en” only).

*isAsync*

An **SLPBoolean** indicating whether the **SLPHandle** should be opened for asynchronous operation or not. AIX supports synchronous operation only.

*phSLP*

A pointer to an **SLPHandle**, in which the open **SLPHandle** is returned. If an error occurs, the value upon return is NULL.

## Return Values

If **SLPOpen** is successful, it returns **SLP\_OK** and an **SLPHandle** handle in the *phSLP* parameter for the language locale passed in as the *pcLang* parameter.

## Error Codes

<b>SLPError</b>	Indicates the status of the operation
-----------------	---------------------------------------

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPParseSrvURL Subroutine,” “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The */etc/slp.conf* File.

The Service Location Protocol (SLP) API.

---

## SLPParseSrvURL Subroutine

### Purpose

Parses the URL passed in as the argument into a service URL structure and returns it in the *ppSrvURL* pointer.

### Syntax

```
SLPError SLPParseSrvURL(char *pcSrvURL
                        SLPSrvURL** ppSrvURL);
```

### Description

The **SLPParseSrvURL** subroutine parses the URL passed in as the argument into a service URL structure and returns it in the *ppSrvURL* pointer. If a parse error occurs, returns **SLP\_PARSE\_ERROR**. The input buffer *pcSrvURL* is destructively modified during the parse and used to fill in the fields of the return structure. The structure returned in *ppSrvURL* should be freed with **SLPFreeURL()**. If the URL has no service part, the **s\_pcSrvPart** string is the empty string (""), not NULL. If *pcSrvURL* is not a service: URL, then the **s\_pcSrvType** field in the returned data structure is the URL's scheme, which might not be the same as the service type under which the URL was registered. If the transport is IP, the **s\_pcTransport** field is the empty string. If the transport is not IP or there is no port number, the **s\_iPort** field is 0.

### Parameters

<i>pcSrvURL</i>	A pointer to a character buffer containing the null-terminated URL string to parse. It is destructively modified to produce the output structure.
<i>ppSrvURL</i>	A pointer to a pointer for the <b>SLPSrvURL</b> structure to receive the parsed URL. The memory should be freed by a call to <b>SLPFree()</b> when no longer needed.

## Return Values

If no error occurs, the return value is **SLP\_OK**. Otherwise, the appropriate error code is returned.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The /etc/slp.conf File.

The Service Location Protocol (SLP) API.

---

## SLPReg Subroutine

### Purpose

Registers the services on the language-specific **SLPHandle** handle and returns the results through the callback.

### Syntax

```
SLPError SLPReg (hSLP, pcSrvURL, usLifetime, pcSrvType, pcAttrs, fresh, callback, pvCookie)
SLPHandle hSLP;
const char *pcSrvURL;
const unsigned short usLifetime;
const char *pcSrvType;
const char *pcAttrs;
SLPBoolean fresh;
SLPRegReport callback;
void *pvCookie;
```

### Description

The **SLPReg** subroutine registers the URL specified by the *pcSrvURL* parameter having the *usLifeTime* lifetime with the attribute list specified by the *pcAttrs* parameter. The attribute list is a comma-separated list of attributes. The *pcSrvType* parameter is the service type name and can be included in the **scheme** service URL that are not in the service. In the case of the **scheme** service URL with service, the *pcSrvType* parameter is ignored. The *fresh* flag specifies that this registration is a new or an update-only registration. If the *fresh* parameter is set to SLP\_TRUE, the registration replaces existing registrations. If the *fresh* parameter is set to SLP\_FALSE, the registration only updates existing registrations. The *usLifeTime* parameter must be nonzero and less than or equal to SLP\_LIFETIME\_MAXIMUM. The registration takes place in the language locale of **hhSLP** handle.

### Parameters

<i>hSLP</i>	The language-specific <b>SLPHandle</b> handle on which to register the services.
<i>pcSrvURL</i>	The URL that needs to be registered.
<i>usLifetime</i>	The time after which the registered URL will expire.
<i>pcSrvType</i>	Specifies the service type name that can be included in the service URL, which is not in the <b>scheme</b> service.
<i>pcAttrs</i>	The comma-separated list of attributes to be registered along with the service URL.
<i>fresh</i>	If the <i>fresh</i> parameter is set to SLP_TRUE, the registration is new; if the <i>fresh</i> parameter is set to SLP_FALSE, this registration updates an existing registration.
<i>callback</i>	A callback function through which the results of the operation are reported.
<i>pvCookie</i>	The memory passed to callback code from the client. The parameter can be set to NULL.

## Return Values

<b>SLP_OK</b>	The subroutine has run successfully.
<b>SLPError</b>	An error occurred.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240, “SLPDereg Subroutine” on page 228, “SLPRegReport Callback Subroutine.”

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) APIs.

---

## SLPRegReport Callback Subroutine

### Description

Returns the same callback type as the **SLPReg** and **SLPDereg** subroutines.

### Syntax

```
typedef void SLPRegReport (hSLP, errCode, pvCookie)
SLPHandle hSLP;
SLPError errCode;
void *pvCookie;
```

### Description

The **SLPSrvURLCallback** type is the type of the callback subroutine parameter to the **SLPFindSrvs** subroutine.

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> handle used to initiate the operation.
<i>errCode</i>	An error code indicating that an error occurred during the operation. The callback must check this error code before processing the parameters. If the error code is not SLP_OK, the API library can choose to terminate the outstanding operation.
<i>pvCookie</i>	The memory passed down from the client code that calls the original API function at the start of the operation. The parameter can be set to NULL.

## Return Values

<b>SLP_TRUE</b>	More data is necessary.
<b>SLP_FALSE</b>	No additional data is necessary.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree

Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine,” “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240, “SLPDereg Subroutine” on page 228, “SLPReg Subroutine” on page 237.

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) APIs.

---

## SLPSrvTypeCallback Subroutine

### Purpose

Returns the same callback type as the **SLPFindSrvTypes()** function.

### Syntax

```
typedef SLPBoolean SLPsrvTypeCallback(SLPHandle hSLP,  
                                       const char* pcSrvTypes,  
                                       SLPErrCode errCode,  
                                       void *pvCookie);
```

### Description

The **SLPSrvTypeCallback** type is the type of the callback function parameter to the **SLPFindSrvTypes()** function.

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> used to initiate the operation.
<i>pcSrvTypes</i>	A character buffer containing a comma-separated, null-terminated list of service types.
<i>errCode</i>	An error code indicating if an error occurred during the operation. The callback should check this error code before processing the parameters. If the error code is other than <b>SLP_OK</b> , then the API library can choose to terminate the outstanding operation.
<i>pvCookie</i>	Memory passed down from the client code that called the original API function, starting the operation. Can be NULL.

### Return Values

The client code should return **SLP\_TRUE** if more data is desired; otherwise **SLP\_FALSE** is returned.

### Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvURLCallback Subroutine” on page 240, “SLPUnescape Subroutine” on page 240

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## SLPSrvURLCallback Subroutine

### Purpose

Returns the same callback type as the **SLPFindSrvs()** function.

### Syntax

```
typedef SLPBoolean SLPsrvURLCallback(SLPHandle hSLP,  
                                     const char* pcSrvURL,  
                                     unsigned short sLifetime,  
                                     SLPError errCode,  
                                     void *pvCookie);
```

### Description

The **SLPSrvURLCallback** type is the type of the callback function parameter to the **SLPFindSrvs()** function.

### Parameters

<i>hSLP</i>	The <b>SLPHandle</b> used to initiate the operation.
<i>pcSrvURL</i>	A character buffer containing the returned service URL.
<i>sLifetime</i>	An unsigned short giving the lifetime of the service advertisement, in seconds. The value must be an unsigned integer less than or equal to <b>SLP_LIFETIME_MAXIMUM</b> .
<i>errCode</i>	An error code indicating if an error occurred during the operation. The callback should check this error code before processing the parameters. If the error code is other than <b>SLP_OK</b> , then the API library can choose to terminate the outstanding operation.
<i>pvCookie</i>	Memory passed down from the client code that called the original API function, starting the operation. Can be NULL.

### Return Values

The client code should return **SLP\_TRUE** if more data is desired; otherwise **SLP\_FALSE** is returned.

### Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPUnescape Subroutine”

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## SLPUnescape Subroutine

### Purpose

Processes an input string and unescapes any characters reserved for SLP.

## Syntax

```
SLPError SLPUnescape(const char* pcInbuf,
                    char** ppcOutBuf,
                    SLPBoolean isTag);
```

## Description

The **SLPUnescape** subroutine processes the input string in *pcInbuf* and unescapes any characters reserved for SLP. If the *isTag* parameter is **SLPTrue**, **SLPUnescape** looks for bad tag characters and signals an error if any are found by returning the **SLP\_PARSE\_ERROR** code. No transformation is performed if the input string is opaque. The results are put into a buffer allocated by the API library and returned in the *ppcOutBuf* parameter. This buffer should be deallocated using **SLPFree()** when the memory is no longer needed.

## Parameters

<i>pcInbuf</i>	Pointer to the input buffer to process for escape characters.
<i>ppcOutBuf</i>	Pointer to a pointer for the output buffer with the characters reserved for SLP escaped. Must be freed using <b>SLPFree()</b> when the memory is no longer needed.
<i>isTag</i>	When true, the input buffer is checked for bad tag characters.

## Return Values

The **SLPUnescape** subroutine returns **SLP\_PARSE\_ERROR** if any characters are bad tag characters and the *isTag* flag is true; otherwise, it returns **SLP\_OK**, or the appropriate error code if another error occurs.

## Related Information

“SLPAttrCallback Subroutine” on page 227, “SLPClose Subroutine” on page 228, “SLPEscape Subroutine” on page 229, “SLPFindAttrs Subroutine” on page 230, “SLPFindScopes Subroutine” on page 231, “SLPFindSrvs Subroutine” on page 232, “SLPFindSrvTypes Subroutine” on page 233, “SLPFree Subroutine” on page 234, “SLPGetProperty Subroutine” on page 234, “SLPOpen Subroutine” on page 235, “SLPParseSrvURL Subroutine” on page 236, “SLPSrvTypeCallback Subroutine” on page 239, “SLPSrvURLCallback Subroutine” on page 240

The `/etc/slp.conf` File.

The Service Location Protocol (SLP) API.

---

## socket Subroutine

### Purpose

Creates an end point for communication and returns a descriptor.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <sys/types.h>
#include <sys/socket.h>
#include <sys/socketvar.h>
int socket ( AddressFamily, Type, Protocol)
int AddressFamily, Type, Protocol;
```

## Description

The **socket** subroutine creates a socket in the specified *AddressFamily* and of the specified type. A protocol can be specified or assigned by the system. If the protocol is left unspecified (a value of 0), the system selects an appropriate protocol from those protocols in the address family that can be used to support the requested socket type.

The **socket** subroutine returns a descriptor (an integer) that can be used in later subroutines that operate on sockets.

Socket level options control socket operations. The **getsockopt** and **setsockopt** subroutines are used to get and set these options, which are defined in the */usr/include/sys/socket.h* file.

## Parameters

<i>AddressFamily</i>	<p>Specifies an address family with which addresses specified in later socket operations should be interpreted. The <i>/usr/include/sys/socket.h</i> file contains the definitions of the address families. Commonly used families are:</p> <p><b>AF_UNIX</b> Denotes the operating system path names.</p> <p><b>AF_INET</b> Denotes the ARPA Internet addresses.</p> <p><b>AF_NS</b> Denotes the XEROX Network Systems protocol.</p> <p><b>AF_BYPASS</b> Denotes the kernel-bypass protocol domain (for example, the protocols that operate over InfiniBand).</p>
<i>Type</i>	<p>Specifies the semantics of communication. The <i>/usr/include/sys/socket.h</i> file defines the socket types. The operating system supports the following types:</p> <p><b>SOCK_STREAM</b> Provides sequenced, two-way byte streams with a transmission mechanism for out-of-band data.</p> <p><b>SOCK_DGRAM</b> Provides datagrams, which are connectionless messages of a fixed maximum length (usually short).</p> <p><b>SOCK_RAW</b> Provides access to internal network protocols and interfaces. This type of socket is available only to the root user, or to non-root users who have the CAP_NUMA_ATTACH capability. (For non-root raw socket access, the CAP_NUMA_ATTACH capability, along with CAP_PROPAGATE, is assigned using the <b>chuser</b> command. For more information about the <b>chuser</b> command, see <b>chuser Command</b> in <i>AIX Version 6.1 Commands Reference, Volume 1</i>.)</p> <p><b>SOCK_SEQPACKET</b> Provides sequenced, reliable, and unduplicated flow of information. This type of socket is used for UDP-style socket creation in case of Stream Control Transmission Protocol and Reliable Datagram Sockets (RDS) Protocol.</p>
<i>Protocol</i>	<p>Specifies a particular protocol to be used with the socket. Specifying the <i>Protocol</i> parameter of 0 causes the <b>socket</b> subroutine to default to the typical protocol for the requested type of returned socket. For SCTP sockets, the protocol parameter is IPPROTO_SCTP. For RDS sockets, the <i>Protocol</i> parameter is BYPASSPROTO_RDS.</p>

## Return Values

Upon successful completion, the **socket** subroutine returns an integer (the socket descriptor).

If the **socket** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*.

## Error Codes

The **socket** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EMFILE</b>	The per-process descriptor table is full.
<b>ENOBUFS</b>	Insufficient resources were available in the system to complete the call.
<b>ESOCKTNOSUPPORT</b>	The socket in the specified address family is not supported.

## Examples

The following program fragment illustrates the use of the **socket** subroutine to create a datagram socket for on-machine use:

```
s = socket(AF_UNIX, SOCK_DGRAM, 0);
```

## Implementation Specifics

The **socket** subroutine is part of Base Operating System (BOS) Runtime.

The **socket** applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Related Information

The **accept** subroutine, **bind** subroutine, **connect** subroutine, **getsockname** subroutine, **getsockopt** subroutine, **ioctl** subroutine, **listen** subroutine, **recv** subroutine, **recvfrom** subroutine, **recvmsg** subroutine, **select** subroutine, **send** subroutine, **sendmsg** subroutine, **sendto** subroutine, **setsockopt** subroutine, **shutdown** subroutine, **socketpair** subroutine.

Initiating Internet Stream Connections Example Program, Sockets Overview, Understanding Socket Creation in *AIX Version 6.1 Communications Programming Concepts*.

---

## socketpair Subroutine

### Purpose

Creates a pair of connected sockets.

### Library

Standard C Library (**libc.a**)

### Syntax

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

```
int socketpair (Domain, Type, Protocol, SocketVector[0])
int Domain, Type, Protocol;
int SocketVector[2];
```

## Description

The **socketpair** subroutine creates an unnamed pair of connected sockets in a specified domain, of a specified type, and using the optionally specified protocol. The two sockets are identical.

**Note:** Create sockets with this subroutine only in the **AF\_UNIX** protocol family.

The descriptors used in referencing the new sockets are returned in the *SocketVector*[0] and *SocketVector*[1] parameters.

The */usr/include/sys/socket.h* file contains the definitions for socket domains, types, and protocols.

All applications containing the **socketpair** subroutine must be compiled with the **\_BSD** macro set to a value of 43 or 44. Socket applications must include the BSD **libbsd.a** library.

## Parameters

<i>Domain</i>	Specifies the communications domain within which the sockets are created. This subroutine does not create sockets in the Internet domain.
<i>Type</i>	Specifies the communications method, whether <b>SOCK_DGRAM</b> or <b>SOCK_STREAM</b> , that the socket uses.
<i>Protocol</i>	Points to an optional identifier used to specify which standard set of rules (such as UDP/IP and TCP/IP) governs the transfer of data.
<i>SocketVector</i>	Points to a two-element vector that contains the integer descriptors of a pair of created sockets.

## Return Values

Upon successful completion, the **socketpair** subroutine returns a value of 0.

If the **socketpair** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable.

## Error Codes

If the **socketpair** subroutine is unsuccessful, it returns one of the following errors codes:

Error	Description
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EFAULT</b>	The <i>SocketVector</i> parameter is not in a writable part of the user address space.
<b>EMFILE</b>	This process has too many descriptors in use.
<b>ENFILE</b>	The maximum number of files allowed are currently open.
<b>ENOBUFS</b>	Insufficient resources were available in the system to perform the operation.
<b>EOPNOTSUPP</b>	The specified protocol does not allow the creation of socket pairs.
<b>EPROTONOSUPPORT</b>	The specified protocol cannot be used on this system.
<b>EPROTOYPE</b>	The socket type is not supported by the protocol.

## Related Information

The **socket** subroutine.

Socketpair Communication Example Program, Sockets Overview, and Understanding Socket Creation in *AIX Version 6.1 Communications Programming Concepts*.

---

## socks5\_getserv Subroutine

### Purpose

Return the address of the SOCKSv5 server (if any) to use when connecting to a given destination.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>

struct sockaddr * socks5_getserv (Dst, DstLen)
struct sockaddr *Dst;
size_t DstLen;
```

### Description

The **socks5\_getserv** subroutine determines which (if any) SOCKSv5 server should be used as an intermediary when connecting to the address specified in *Dst*.

The address returned in *Dst* may be IPv4 or IPv6 or some other family. The user should check the address family before using the returned data.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

### Parameters

<i>Dst</i>	Specifies the external address of the target socket to use as a key for looking up the appropriate SOCKSv5 server.
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .

### Return Values

- Upon successful lookup, the **socks\_getserv** subroutine returns a reference to a **sockaddr** struct.
- If the **socks5tcp\_connect** subroutine is unsuccessful in finding a server, for any reason, a value of NULL is returned. If an error occurred, an error code, indicating the generic system error, is moved into the **errno** global variable.

### Error Codes (placed in errno)

The **socks5\_getserv** subroutine is unsuccessful if no server is indicated or if any of the following errors occurs:

Error	Description
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EFAULT</b>	The <i>Dst</i> parameter is not in a writable part of the user address space.
<b>EINVAL</b>	One or more of the specified arguments is invalid.
<b>ENOMEM</b>	The <i>Dst</i> parameter is not large enough to hold the server address.

### Examples

The following program fragment illustrates the use of the **socks5\_getserv** subroutine by a client to request a connection from a server's socket.

```

struct sockaddr_in6 dst;

struct sockaddr *srv;
.
.
.
srv = socks5_getserv((struct sockaddr*)&dst, sizeof(dst));

if (srv !=NULL) {

    /* Success: srv should be used as the socks5 server */

} else {
    /* Failure: no server could be returned.  check errno */
}

```

## Related Information

The **socks5tcp\_connect** subroutine, **socks5tcp\_bind** subroutine, **socks5tcp\_accept** subroutine, **socks5udp\_associate** subroutine, **socks5udp\_sendto** subroutine, **/etc/socks5c.conf** file, **connect** subroutine.

Sockets Overview and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

SOCKS5C\_CONFIG Environment Variable in *AIX Version 6.1 Files Reference*.

---

## /etc/socks5c.conf File

### Purpose

Contains mappings between network destinations and SOCKSv5 servers.

### Description

The **/etc/socks5c.conf** file contains basic mappings between network destinations (hosts or networks) and SOCKSv5 servers to use when accessing those destinations. This is an ASCII file that contains records for server mappings. Text following a pound character ('#') is ignored until the end of line. Each record appears on a single line and is the following format:

```
<destination>[/<prefixlength>]  <server>[:<port>]
```

You must separate fields with whitespace. Records are separated by new-line characters. The fields and modifiers in a record have the following values:

<i>destination</i>	Specifies a network destination; <i>destination</i> may be either a name fragment or a numeric address (with optional <i>prefixlength</i> ). If <i>destination</i> is an address, it may be either IPv4 or IPv6.
<i>prefixlength</i>	If specified, indicates the number of leftmost (network order) bits of an address to use when comparing to this record. Only valid if <i>destination</i> is an address. If not specified, all bits are used in comparisons.
<i>server</i>	Specifies the SOCKSv5 server associated with <i>destination</i> . If <i>server</i> is "NONE" (must be all uppercase), this record indicates that target addresses matching <i>destination</i> should not use any SOCKSv5 server, but rather be contacted directly.
<i>port</i>	If specified, indicates the port to use when contacting <i>server</i> . If not specified, the default of 1080 is assumed. <b>Note:</b> Server address in IPv6 format <b>must</b> be followed by a port number.

If a name fragment *destination* is present in **/etc/socks5c.conf**, all target addresses is SOCKSv5 operations will be converted into hostnames for name comparison (in addition to numeric comparisons with numeric records). The resulting hostname is considered to match if the last characters in the hostname match the specified name fragment.

When using this configuration information to determine the address of the appropriate SOCKSv5 server for a target destination, the "best" match is used. The "best" match is defined as:

<i>destination</i> is numeric	Most bits in comparison (i.e. largest <i>prefixlength</i> )
<i>destination</i> is a name fragment	Most characters in name fragment.

When both name fragment and numeric addresses are present, all name fragment entries are "better" than numeric address entries.

Two implicit records:

```
0.0.0.0/0 NONE #All IPv4 destinations; no associated server.
```

```
::/0 NONE #All IPv6 destinations; no associated server.
```

are assumed as defaults for all destinations not specified in **/etc/socks5c.conf**.

## Security

Access Control: This file should grant read (r) access to all users and grant write (w) access only to the root user.

## Examples

```
#Sample socks5c.conf file
```

```
9.0.0.0/8 NONE #Direct communication with all hosts in the 9 network.
```

```
129.35.0.0/16 sox1.austin.ibm.com
```

```
ibm.com NONE #Direct communication will all hosts matching "ibm.com" (e.g. "aguila.austin.ibm.com")
```

## Related Information

The **socks5tcp\_connect** subroutine, **socks5tcp\_bind** subroutine, **socks5tcp\_accept** subroutine, **socks5udp\_associate** subroutine, **socks5udp\_sendto** subroutine, **socks5\_getserv** subroutine, **connect** subroutine.

---

## socks5tcp\_accept Subroutine

### Purpose

Awaits an incoming connection to a socket from a previous **socks5tcp\_bind()** call.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>
```

```

int socks5tcp_accept(Socket, Dst, DstLen, Svr, SvrLen)
int Socket;
struct sockaddr *Dst;
size_t DstLen;
struct sockaddr *Svr;
size_t SvrLen;

```

## Description

The **socks5tcp\_accept** subroutine blocks until an incoming connection is established on a listening socket that was requested in a previous call to **socks5tcp\_bind**. Upon success, subsequent writes to and reads from *Socket* will be relayed through *Svr*.

*Socket* must be an open socket descriptor of type SOCK\_STREAM.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Dst</i>	If non-NULL, buffer for receiving the address of the remote client which initiated an incoming connection
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .
<i>Svr</i>	If non-NULL, specifies the address of the SOCKSv5 server to use to request the relayed connection; on success, this space will be overwritten with the server-side address of the incoming connection.
<i>SvrLength</i>	Specifies the length of the address structure in <i>Svr</i> .

## Return Values

Upon successful completion, the **socks5tcp\_accept** subroutine returns a value of 0, and modifies *Dst* and *Svr* to reflect the actual endpoints of the incoming external socket.

If the **socks5tcp\_accept** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the generic system error, into the **errno** global variable.
- Moves an error code, indicating the specific SOCKSv5 error, into the **socks5\_errno** global variable.

## Error Codes (placed in errno; inherited from underlying call to connect())

The **socks5tcp\_bindaccept** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>EFAULT</b>	The <i>Dst</i> or <i>Svr</i> parameter is not in a writable part of the user address space.
<b>EINVAL</b>	One or more of the specified arguments is invalid.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENOSPC</b>	There is no space left on a device or system table.
<b>ENOTCONN</b>	The socket could not be connected.

## Error Codes (placed in `socks5_errno`; SOCKSv5-specific errors)

The `socks5tcp_connect` subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<code>S5_ESRVFAIL</code>	General SOCKSv5 server failure.
<code>S5_EPERM</code>	SOCKSv5 server ruleset rejection.
<code>S5_ENETUNREACH</code>	SOCKSv5 server could not reach target network.
<code>S5_EHOSTUNREACH</code>	SOCKSv5 server could not reach target host.
<code>S5_ECONNREFUSED</code>	SOCKSv5 server connection request refused by target host.
<code>S5_ETIMEDOUT</code>	SOCKSv5 server connection failure due to TTL expiry.
<code>S5_EOPNOTSUPP</code>	Command not supported by SOCKSv5 server.
<code>S5_EAFNOSUPPORT</code>	Address family not supported by SOCKSv5 server.
<code>S5_EADDRINUSE</code>	Requested bind address is already in use (at the SOCKSv5 server).
<code>S5_ENOSERV</code>	No server found.

## Examples

The following program fragment illustrates the use of the `socks5tcp_accept` and `socks5tcp_bind` subroutines by a client to request a listening socket from a server and wait for an incoming connection on the server side.

```
struct sockaddr_in svr;
struct sockaddr_in dst;
.
.
.
socks5tcp_bind(s, (struct sockaddr*)&dst, sizeof(dst), (struct sockaddr *)&svr, sizeof(svr), &res, sizeof(svr));
.
.
.
socks5tcp_accept(s, (struct sockaddr *)&dst, sizeof(dst), (struct sockaddr *)&svr, sizeof(svr));
```

## Related Information

The `socks5tcp_connect` subroutine, `socks5tcp_bind` subroutine, `socks5udp_associate` subroutine, `socks5udp_sendto` subroutine, `socks5_getserv` subroutine, `/etc/socks5c.conf` file, `accept` subroutine, `bind` subroutine, `getsockname` subroutine, `send` subroutine, `socket` subroutine.

Initiating UNIX Stream Connections Example Program, Sockets Overview, and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

SOCKS5C\_CONFIG Environment Variable in *AIX Version 6.1 Files Reference*.

---

## socks5tcp\_bind Subroutine

### Purpose

Connect to a SOCKSv5 server and request a listening socket for incoming remote connections.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>
```

```
int socks5tcp_bind(Socket, Dst, DstLen, Svr, SvrLen)
Int Socket;
struct sockaddr *Dst;
size_t DstLen;
struct sockaddr *Svr;
size_t SvrLen;
```

## Description

The **socks5tcp\_bind** subroutine requests a listening socket on the SOCKSv5 server specified in *Svr*, in preparation for an incoming connection from a remote destination, specified by *Dst*. Upon success, *Svr* will be overwritten with the actual address of the newly bound listening socket, and *Socket* may be used in a subsequent call to **socks5tcp\_accept**.

*Socket* must be an open socket descriptor of type SOCK\_STREAM.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Dst</i>	Specifies the address of the SOCKSv5 server to use to request the relayed connection; on success, this space will be overwritten with the actual bound address on the server.
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .
<i>Svr</i>	If non-NULL, specifies the address of the SOCKSv5 server to use to request the relayed connection; on success, this space will be overwritten with the actual bound address on the server.
<i>SvrLength</i>	Specifies the length of the address structure in <i>Svr</i> .

## Return Values

Upon successful completion, the **socks5tcp\_bind** subroutine returns a value of 0, and modifies *Svr* to reflect the actual address of the newly bound listener socket.

If the **socks5tcp\_bind** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the generic system error, into the **errno** global variable.
- Moves an error code, indicating the specific SOCKSv5 error, into the **socks5\_errno** global variable.

## Error Codes (placed in errno; inherited from underlying call to connect())

The **socks5tcp\_bindaccept** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EISCONN</b>	The socket is already connected.
<b>ETIMEDOUT</b>	The establishment of a connection timed out before a connection was made.
<b>ECONNREFUSED</b>	The attempt to connect was rejected.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>EADDRINUSE</b>	The specified address is already in use.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINPROGRESS</b>	The socket is marked as nonblocking. The connection cannot be immediately completed. The application program can select the socket for writing during the connection process.

Error	Description
<b>EINVAL</b>	One or more of the specified arguments is invalid.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENOSPC</b>	There is no space left on a device or system table.
<b>ENOTCONN</b>	The socket could not be connected.

## Error Codes (placed in `socks5_errno`; SOCKSv5-specific errors)

The **socks5tcp\_connect** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>S5_ESRVFAIL</b>	General SOCKSv5 server failure.
<b>S5_EPERM</b>	SOCKSv5 server ruleset rejection.
<b>S5_ENETUNREACH</b>	SOCKSv5 server could not reach target network.
<b>S5_EHOSTUNREACH</b>	SOCKSv5 server could not reach target host.
<b>S5_ECONNREFUSED</b>	SOCKSv5 server connection request refused by target host.
<b>S5_ETIMEDOUT</b>	SOCKSv5 server connection failure due to TTL expiry.
<b>S5_EOPNOTSUPP</b>	Command not supported by SOCKSv5 server.
<b>S5_EAFNOSUPPORT</b>	Address family not supported by SOCKSv5 server.
<b>S5_EADDRINUSE</b>	Requested bind address is already in use (at the SOCKSv5 server).
<b>S5_ENOERV</b>	No server found.

## Examples

The following program fragment illustrates the use of the **socks5tcp\_bind** subroutine by a client to request a listening socket from a server.

```
struct sockaddr_in svr;
struct sockaddr_in dst;
.
.
.
socks5tcp_bind(s, (struct sockaddr *)&dst, sizeof(dst), (structsockaddr *)&svr, sizeof(svr));
```

## Related Information

The **socks5tcp\_accept** subroutine, **socks5tcp\_connect** subroutine, **socks5\_getserv** subroutine, **/etc/socks5c.conf** file, **accept** subroutine, **bind** subroutine, **getsockname** subroutine, **send** subroutine, **socket** subroutine.

Sockets Overview and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

SOCKS5C\_CONFIG Environment Variable in *AIX Version 6.1 Files Reference*.

---

## socks5tcp\_connect Subroutine

### Purpose

Connect to a SOCKSv5 server and request a connection to an external destination.

### Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>

int socks5tcp_connect (Socket, Dst, DstLen, Svr, SvrLen)
int Socket;
struct sockaddr *Dst;
size_t DstLen;
struct sockaddr *Svr;
size_t SvrLen;
```

## Description

The **socks5tcp\_connect** subroutine requests a connection to *Dst* from the SOCKSv5 server specified in *Svr*. If successful, *Dst* and *Svr* will be overwritten with the actual addresses of the external connection and subsequent writes to and reads from *Socket* will be relayed through *Svr*.

*Socket* must be an open socket descriptor of type SOCK\_STREAM; *Dst* and *Svr* may be either IPv4 or IPv6 addresses.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Dst</i>	Specifies the external address of the target socket to which the SOCKSv5 server will attempt to connect.
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .
<i>Svr</i>	If non-NULL, specifies the address of the SOCKSv5 server to use to request the relayed connection.
<i>SvrLength</i>	Specifies the length of the address structure in <i>Svr</i> .

## Return Values

Upon successful completion, the **socks5tcp\_connect** subroutine returns a value of 0, and modifies *Dst* and *Svr* to reflect the actual endpoints of the created external socket.

If the **socks5tcp\_connect** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the generic system error, into the **errno** global variable.
- Moves an error code, indicating the specific SOCKSv5 error, into the **socks5\_errno** global variable.
- *Dst* and *Svr* are left unmodified.

## Error Codes (placed in errno; inherited from underlying call to connect())

The **socks5tcp\_connect** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EADDRNOTAVAIL</b>	The specified address is not available from the local machine.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>EISCONN</b>	The socket is already connected.

Error	Description
<b>ETIMEDOUT</b>	The establishment of a connection timed out before a connection was made.
<b>ECONNREFUSED</b>	The attempt to connect was rejected.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>EADDRINUSE</b>	The specified address is already in use.
<b>EFAULT</b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b>EINPROGRESS</b>	The socket is marked as nonblocking. The connection cannot be immediately completed. The application program can select the socket for writing during the connection process.
<b>EINVAL</b>	One or more of the specified arguments is invalid.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENOSPC</b>	There is no space left on a device or system table.
<b>ENOTCONN</b>	The socket could not be connected.

## Error Codes (placed in `socks5_errno`; SOCKSv5-specific errors)

The `socks5tcp_connect` subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>S5_ESRVFAIL</b>	General SOCKSv5 server failure.
<b>S5_EPERM</b>	SOCKSv5 server ruleset rejection.
<b>S5_ENETUNREACH</b>	SOCKSv5 server could not reach target network.
<b>S5_EHOSTUNREACH</b>	SOCKSv5 server could not reach target host.
<b>S5_ECONNREFUSED</b>	SOCKSv5 server connection request refused by target host.
<b>S5_ETIMEDOUT</b>	SOCKSv5 server connection failure due to TTL expiry.
<b>S5_EOPNOTSUPP</b>	Command not supported by SOCKSv5 server.
<b>S5_EAFNOSUPPORT</b>	Address family not supported by SOCKSv5 server.
<b>S5_ENOSERV</b>	No server found.

## Examples

The following program fragment illustrates the use of the `socks5tcp_connect` subroutine by a client to request a connection from a server's socket.

```
struct sockaddr_in svr;
struct sockaddr_in6 dst;
.
.
.
socks5tcp_connect(s, (struct sockaddr*)&dst, sizeof(dst), (struct sockaddr *)&svr, sizeof(svr));
```

## Related Information

The `socks5_getserv` subroutine, `/etc/socks5c.conf` file, `connect` subroutine, `accept` subroutine, `bind` subroutine, `getsockname` subroutine, `send` subroutine, `socket` subroutine.

Initiating UNIX Stream Connections Example Program, Sockets Overview, and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

SOCKS5C\_CONFIG Environment Variable in *AIX Version 6.1 Files Reference*.

---

## socks5udp\_associate Subroutine

### Purpose

Connects to a SOCKSv5 server, and requests a UDP association for subsequent UDP socket communications.

## Library

Standard C Library (**libc.a**)

## Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>

int socks5udp_associate (Socket, Dst, DstLen, Svr, SvrLen)
int Socket;
const struct sockaddr *Dst;
size_t DstLen;
const struct sockaddr *Svr;
size_t SvrLen;
```

## Description

The **socks5udp\_associate** subroutine requests a UDP association for *Dst* on the SOCKSv5 server specified in *Svr*. Upon success, *Dst* is overwritten with a rendezvous address to which subsequent UDP packets should be sent for relay by *Svr*.

*Socket* must be an open socket descriptor of type SOCK\_STREAM; *Dst* and *Svr* may be either IPv4 or IPv6 addresses.

Note that *Socket* cannot be used to send subsequent UDP packets (a second socket of type SOCK\_DGRAM must be created).

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

## Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Dst</i>	Specifies the external address of the target socket to which the SOCKSv5 client expects to send UDP packets.
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .
<i>Svr</i>	Specifies the address of the SOCKSv5 server to use to request the association.
<i>SvrLength</i>	Specifies the length of the address structure in <i>Svr</i> .

## Return Values

Upon successful completion, the **socks5udp\_associate** subroutine returns a value of 0 and overwrites *Dst* with the rendezvous address.

If the **socks5udp\_associate** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the generic system error, into the **errno** global variable.
- Moves an error code, indicating the specific SOCKSv5 error, into the **socks5\_errno** global variable.

## Error Codes (placed in `errno`; inherited from underlying call to `connect()`)

The **`socks5udp_associate`** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b><code>EBADF</code></b>	The <i>Socket</i> parameter is not valid.
<b><code>ENOTSOCK</code></b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b><code>EADDRNOTAVAIL</code></b>	The specified address is not available from the local machine.
<b><code>EAFNOSUPPORT</code></b>	The addresses in the specified address family cannot be used with this socket.
<b><code>EISCONN</code></b>	The socket is already connected.
<b><code>ETIMEDOUT</code></b>	The establishment of a connection timed out before a connection was made.
<b><code>ECONNREFUSED</code></b>	The attempt to connect was rejected.
<b><code>ENETUNREACH</code></b>	No route to the network or host is present.
<b><code>EADDRINUSE</code></b>	The specified address is already in use.
<b><code>EFAULT</code></b>	The <i>Address</i> parameter is not in a writable part of the user address space.
<b><code>EINPROGRESS</code></b>	The socket is marked as nonblocking. The connection cannot be immediately completed. The application program can select the socket for writing during the connection process.
<b><code>EINVAL</code></b>	One or more of the specified arguments is invalid.
<b><code>ENETDOWN</code></b>	The specified physical network is down.
<b><code>ENOSPC</code></b>	There is no space left on a device or system table.
<b><code>ENOTCONN</code></b>	The socket could not be connected.

## Error Codes (placed in `socks5_errno`; SOCKSv5-specific errors)

The **`socks5tcp_connect`** subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b><code>S5_ESRVFAIL</code></b>	General SOCKSv5 server failure.
<b><code>S5_EPERM</code></b>	SOCKSv5 server ruleset rejection.
<b><code>S5_ENETUNREACH</code></b>	SOCKSv5 server could not reach target network.
<b><code>S5_EHOSTUNREACH</code></b>	SOCKSv5 server could not reach target host.
<b><code>S5_ECONNREFUSED</code></b>	SOCKSv5 server connection request refused by target host.
<b><code>S5_ETIMEDOUT</code></b>	SOCKSv5 server connection failure due to TTL expiry.
<b><code>S5_EOPNOTSUPP</code></b>	Command not supported by SOCKSv5 server.
<b><code>S5_EAFNOSUPPORT</code></b>	Address family not supported by SOCKSv5 server.
<b><code>S5_ENOSERV</code></b>	No server found.

## Examples

The following program fragment illustrates the use of the **`socks5udp_associate`** subroutine by a client to request an association on a server.

```
struct sockaddr_in svr;  
struct sockaddr_in6 dst;  
.  
.  
.  
socks5udp_associate(s, (struct sockaddr*)&dst, sizeof(dst), (struct sockaddr*)&svr, sizeof(svr));
```

## Related Information

The **`socks5udp_sendto`** subroutine, **`socks5_getserv`** subroutine, **`/etc/socks5c.conf`** file, **`connect`** subroutine, **`accept`** subroutine, **`bind`** subroutine, **`getsockname`** subroutine, **`send`** subroutine, **`socket`** subroutine.

Initiating UNIX Stream Connections Example Program, Sockets Overview, and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

## socks5udp\_sendto Subroutine

### Purpose

Send UDP packets through a SOCKSv5 server.

### Library

Standard C Library (**libc.a**)

### Syntax

```
#include <stdlib.h>
#include <netinet/in.h>
#include <sys/socket.h>

int socks5udp_sendto (Socket, Message, MsgLen, Flags, Dst, DstLen, Svr, SvrLen)
int Socket;
void *Message;
size_t MsgLen;
int Flags;
struct sockaddr *Dst;
size_t DstLen;
struct sockaddr *Svr;
size_t SvrLen;
```

### Description

The **socks5udp\_sendto** subroutine sends a UDP packet to *Svr* for relay to *Dst*. *Svr* must be the rendezvous address returned from a previous call to **socks5udp\_associate**.

*Socket* must be an open socket descriptor of type SOCK\_DGRAM; *Dst* and *Svr* may be either IPv4 or IPv6 addresses.

The socket applications can be compiled with **COMPAT\_43** defined. This will make the **sockaddr** structure BSD 4.3 compatible. For more details refer to **socket.h**.

### Parameters

<i>Socket</i>	Specifies the unique name of the socket.
<i>Message</i>	Specifies the address containing the message to be sent.
<i>MsgLen</i>	Specifies the size of the message in bytes.
<i>Flags</i>	Allows the sender to control the message transmission. See the description in the <b>sendto</b> subroutine for more specific details.
<i>Dst</i>	Specifies the external address to which the SOCKSv5 server will attempt to relay the UDP packet.
<i>DstLength</i>	Specifies the length of the address structure in <i>Dst</i> .
<i>Svr</i>	Specifies the address of the SOCKSv5 server to send the UDP packet for relay.
<i>SvrLength</i>	Specifies the length of the address structure in <i>Svr</i> .

### Return Values

Upon successful completion, the **socks5udp\_sendto** subroutine returns a value of 0.

If the **socks5udp\_sendto** subroutine is unsuccessful, the system handler performs the following functions:

- Returns a value of -1 to the calling program.
- Moves an error code, indicating the generic system error, into the **errno** global variable.
- Moves an error code, indicating the specific SOCKSv5 error, into the **socks5\_errno** global variable.

## Error Codes (placed in `errno`; inherited from underlying call to `sendto()`)

The `socks5tcp_connect` subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>EBADF</b>	The <i>Socket</i> parameter is not valid.
<b>ENOTSOCK</b>	The <i>Socket</i> parameter refers to a file, not a socket.
<b>EAFNOSUPPORT</b>	The addresses in the specified address family cannot be used with this socket.
<b>ENETUNREACH</b>	No route to the network or host is present.
<b>EINVAL</b>	One or more of the specified arguments is invalid.
<b>ENETDOWN</b>	The specified physical network is down.
<b>ENOSPC</b>	There is no space left on a device or system table.

## Error Codes (placed in `socks5_errno`; SOCKSv5-specific errors)

The `socks5tcp_connect` subroutine is unsuccessful if any of the following errors occurs:

Error	Description
<b>S5_ESRVFAIL</b>	General SOCKSv5 server failure.
<b>S5_EPERM</b>	SOCKSv5 server ruleset rejection.
<b>S5_ENETUNREACH</b>	SOCKSv5 server could not reach target network.
<b>S5_EHOSTUNREACH</b>	SOCKSv5 server could not reach target host.
<b>S5_ECONNREFUSED</b>	SOCKSv5 server connection request refused by target host.
<b>S5_ETIMEDOUT</b>	SOCKSv5 server connection failure due to TTL expiry.
<b>S5_EOPNOTSUPP</b>	Command not supported by SOCKSv5 server.
<b>S5_EAFNOSUPPORT</b>	Address family not supported by SOCKSv5 server.
<b>S5_ENOSERV</b>	No server found.

## Examples

The following program fragment illustrates the use of the `socks5udp_sendto` subroutine by a client to request a connection from a server's socket.

```
void *message;
size_t msglen;
int flags;
struct sockaddr_in svr;
struct sockaddr_in6 dst;
.
.
.
socks5udp_associate(s, (struct sockaddr*)&dst, sizeof(dst), (struct sockaddr *)&svr, sizeof(svr));
.
.
.
socks5udp_sendto(s, message, msglen, flags (struct sockaddr*)&dst, sizeof(dst), (struct sockaddr *)&svr, sizeof(svr));
```

## Related Information

The `socks5udp_associate` subroutine, `socks5_getserv` subroutine, `/etc/socks5c.conf` file, `bind` subroutine, `getsockname` subroutine, `sendto` subroutine, `socket` subroutine.

Initiating UNIX Stream Connections Example Program, Sockets Overview, and Understanding Socket Connections in *AIX Version 6.1 Communications Programming Concepts*.

SOCKS5C\_CONFIG Environment Variable in *AIX Version 6.1 Files Reference*.

---

## splice Subroutine

### Purpose

Lets the protocol stack manage two sockets that use TCP.

### Syntax

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

```
int splice(socket1, socket2, flags)
    int socket1, socket2;
    int flags;
```

### Description

The **splice** subroutine will let TCP manage two sockets that are in connected state thus relieving the caller from moving data from one socket to another. After the **splice** subroutine returns successfully, the caller needs to close the two sockets.

The two sockets should be of type **SOCK\_STREAM** and protocol **IPPROTO\_TCP**. Specifying a protocol of zero will also work.

### Parameters

<i>socket1, socket2</i>	Specifies a socket that had gone through a successful connect() or accept().
<i>flags</i>	Set to zero. Currently ignored.

### Return Values

0	Indicates a successful completion.
-1	Indicates an error. The specific error is indicated by errno.

### Error Codes

<b>EBADF</b>	<i>socket1</i> or <i>socket2</i> is not valid.
<b>ENOTSOCK</b>	<i>socket1</i> or <i>socket2</i> refers to a file, not a socket.
<b>EOPNOTSUPP</b>	<i>socket1</i> or <i>socket2</i> is not of type <b>SOCK_STREAM</b> .
<b>EINVAL</b>	The parameters are invalid.
<b>EEXIST</b>	<i>socket1</i> or <i>socket2</i> is already spliced.
<b>ENOTCONN</b>	<i>socket1</i> or <i>socket2</i> is not in connected state.
<b>EAFNOSUPPORT</b>	<i>socket1</i> or <i>socket2</i> address family is not supported for this subroutine.

---

## WriteFile Subroutine

### Purpose

Writes data to a socket.

## Syntax

```
#include <iocp.h>
boolean_t WriteFile (FileDescriptor, Buffer, WriteCount, AmountWritten, Overlapped)
HANDLE FileDescriptor;
LPVOID Buffer;
DWORD WriteCount;
LPDWORD AmountWritten;
LPOVERLAPPED Overlapped;
```

## Description

The **WriteFile** subroutine writes the number of bytes specified by the *WriteCount* parameter from the buffer indicated by the *Buffer* parameter to the *FileDescriptor* parameter. The number of bytes written is saved in the *AmountWritten* parameter. The *Overlapped* parameter indicates whether or not the operation can be handled asynchronously.

The **WriteFile** subroutine returns a boolean (an integer) indicating whether or not the request has been completed.

The **WriteFile** subroutine is part of the I/O Completion Port (IOCP) kernel extension.

**Note:** This subroutine only works to a socket file descriptor. It does not work with files or other file descriptors.

## Parameters

<i>FileDescriptor</i>	Specifies a valid file descriptor obtained from a call to the <b>socket</b> or <b>accept</b> subroutines.
<i>Buffer</i>	Specifies the buffer from which the data will be written.
<i>WriteCount</i>	Specifies the maximum number of bytes to write.
<i>AmountWritten</i>	Specifies the number of bytes written. The parameter is set by the subroutine.
<i>Overlapped</i>	Specifies an overlapped structure indicating whether or not the request can be handled asynchronously.

## Return Values

Upon successful completion, the **WriteFile** subroutine returns a boolean indicating the request has been completed.

If the **WriteFile** subroutine is unsuccessful, the subroutine handler performs the following functions:

- Returns a value of 0 to the calling program.
- Moves an error code, indicating the specific error, into the **errno** global variable. For further explanation of the **errno** variable, see the link in the Related Information section of this document.

## Error Codes

<b>EINPROGRESS</b>	The write request can not be immediately satisfied and will be handled asynchronously. A completion packet will be sent to the associated completion port upon completion.
<b>EAGAIN</b>	The write request cannot be immediately satisfied and cannot be handled asynchronously.
<b>EINVAL</b>	The <i>FileDescriptor</i> is invalid.

## Examples

The following program fragment illustrates the use of the **WriteFile** subroutine to synchronously write data to a socket:

```
void buffer;  
int amount_written;  
b=WriteFile (34, &buffer, 128, &amount_written, NULL);
```

The following program fragment illustrates the use of the **WriteFile** subroutine to asynchronously write data to a socket:

```
void buffer;  
int amount_written;  
LPOVERLAPPED overlapped;  
b = ReadFile (34, &buffer, 128, &amount_written, overlapped);
```

**Note:** The request will only be handled asynchronously if it cannot be immediately satisfied.

## Related Information

The “socket Subroutine” on page 241, “accept Subroutine” on page 29, “CreateIoCompletionPort Subroutine” on page 37, “ReadFile Subroutine” on page 165, “GetQueuedCompletionStatus Subroutine” on page 101, “GetMultipleCompletionStatus Subroutine” on page 83, and “PostQueuedCompletionStatus Subroutine” on page 161.

For further explanation of the **errno** variable, see Error Notification Object Class in *AIX Version 6.1 General Programming Concepts: Writing and Debugging Programs*

---

## Chapter 3. Streams

---

### adjmsg Utility

#### Purpose

Trims bytes in a message.

#### Syntax

```
int adjmsg (mp, len)
mb1k_t * mp;
register int len;
```

#### Description

The **adjmsg** utility trims bytes from either the head or tail of the message specified by the *mp* parameter. It only trims bytes across message blocks of the same type. The **adjmsg** utility is unsuccessful if the *mp* parameter points to a message containing fewer than *len* bytes of similar type at the message position indicated.

This utility is part of STREAMS Kernel Extensions.

#### Parameters

*mp* Specifies the message to be trimmed.  
*len* Specifies the number of bytes to remove from the message.

If the value of the *len* parameter is greater than 0, the **adjmsg** utility removes the number of bytes specified by the *len* parameter from the beginning of the *mp* message. If the value of the *len* parameter is less than 0, it removes *len* bytes from the end of the *mp* message. If the value of the *len* parameter is 0, the **adjmsg** utility does nothing.

#### Return Values

On successful completion, the **adjmsg** utility returns a value of 1. Otherwise, it returns a value of 0.

#### Related Information

The **msgdsize** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

### alloca Utility

#### Purpose

Allocates message and data blocks.

#### Syntax

```
struct msgb *
alloca(size, pri)
register int size;
uint pri;
```

## Description

The **allocb** utility allocates blocks for a message. When a message is allocated in this manner, the `b_band` field of the **mblk\_t** structure is initially set to a value of 0. Modules and drivers can set this field.

This utility is part of STREAMS Kernel Extensions.

## Parameters

- size* Specifies the minimum number of bytes needed in the data buffer.
- pri* Specifies the relative importance of the allocated blocks to the module. The possible values are:
- **BPRI\_LO**
  - **BPRI\_MED**
  - **BPRI\_HI**

## Return Values

The **allocb** utility returns a pointer to a message block of type **M\_DATA** in which the data buffer contains at least the number of bytes specified by the *size* parameter. If a block cannot be allocated as requested, the **allocb** utility returns a null pointer.

## Related Information

The **esballoc** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## backq Utility

### Purpose

Returns a pointer to the queue behind a given queue.

### Syntax

```
queue_t *  
backq(q)  
register queue_t * q;
```

### Description

The **backq** utility returns a pointer to the queue preceding a given queue. If no such queue exists (as when the *q* parameter points to a stream end), the **backq** utility returns a null pointer.

This utility is part of STREAMS Kernel Extensions.

### Parameters

- q* Specifies the queue from which to begin.

### Return Values

The **backq** utility returns a pointer to the queue behind a given queue. If no such queue exists, the **backq** utility returns a null pointer.

## Related Information

The **RD** utility, **WR** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## bcanput Utility

### Purpose

Tests for flow control in the given priority band.

### Syntax

```
int
bcanput(q, pri)
register queue_t * q;
unsigned char pri;
```

### Description

The **bcanput** utility provides modules and drivers with a way to test flow control in the given priority band.

The **bcanput (q, 0)** call is equivalent to the **canput (q)** call.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue from which to begin to test.  
*pri* Specifies the priority band to test.

### Return Values

The **bcanput** utility returns a value of 1 if a message of the specified priority can be placed on the queue. It returns a value of 0 if the priority band is flow-controlled and sets the **QWANTW** flag to 0 band (the **QB\_WANTW** flag is set to nonzero band). If the band does not yet exist on the queue in question, it returns a value of 1.

## Related Information

List of Streams Programming References, Understanding STREAMS Flow Control in *AIX Version 6.1 Communications Programming Concepts*.

---

## bufcall Utility

### Purpose

Recovers from a failure of the **allocb** utility.

### Syntax

```
#include <sys/stream.h>

int
bufcall(size, pri, func, arg)
uint size;
```

```
int  pri;
void (* func)();
long arg;
```

## Description

The **bufcall** utility assists in the event of a block-allocation failure. If the **allocb** utility returns a null, indicating a message block is not currently available, the **bufcall** utility may be invoked.

The **bufcall** utility arranges for *(\*func)(arg)* call to be made when a buffer of the number of bytes specified by the *size* parameter is available. The *pri* parameter is as described in the **allocb** utility. When the function specified by the *func* parameter is called, it has no user context. It cannot reference the **u\_area** and must return without sleeping. The **bufcall** utility does not guarantee that the desired buffer will be available when the function specified by the *func* parameter is called since interrupt processing may acquire it.

On an unsuccessful return, the function specified by the *func* parameter will never be called. A failure indicates a temporary inability to allocate required internal data structures.

On multiprocessor systems, the function specified by the *func* parameter should be interrupt-safe. Otherwise, the **STR\_QSAFETY** flag must be set when installing the module or driver with the **str\_install** utility.

This utility is part of STREAMS Kernel Extensions.

**Note:** The **stream.h** header file must be the last included header file of each source file using the stream library.

## Parameters

<i>size</i>	Specifies the number of bytes needed.
<i>pri</i>	Specifies the relative importance of the allocated blocks to the module. The possible values are: <ul style="list-style-type: none"><li>• <b>BPRI_LO</b></li><li>• <b>BPRI_MED</b></li><li>• <b>BPRI_HI</b></li></ul>
<i>func</i>	Specifies the function to be called.
<i>arg</i>	Specifies an argument passed to the function.

## Return Values

The **bufcall** utility returns a value of 1 when the request is successfully recorded. Otherwise, it returns a value of 0.

## Related Information

The **allocb** utility, **unbufcall** utility, **mi\_bufcall** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

Understanding STREAMS Synchronization in *AIX Version 6.1 Communications Programming Concepts*.

---

## canput Utility

### Purpose

Tests for available room in a queue.

### Syntax

```
int  
canput(q)  
register queue_t * q;
```

### Description

The **canput** utility determines if there is room left in a message queue. If the queue does not have a service procedure, the **canput** utility searches farther in the same direction in the stream until it finds a queue containing a service procedure. This is the first queue on which the passed message can actually be queued. If such a queue cannot be found, the search terminates on the queue at the end of the stream.

The **canput** utility only takes into account normal data flow control.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue at which to begin the search.

### Return Values

The **canput** utility tests the queue found by the search. If the message queue in this queue is not full, the **canput** utility returns a value of 1. This return indicates that a message can be put to the queue. If the message queue is full, the **canput** utility returns a value of 0. In this case, the caller is generally referred to as "blocked".

### Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## clone Device Driver

### Purpose

Opens an unused minor device on another STREAMS driver.

### Description

The **clone** device driver is a STREAMS software driver that finds and opens an unused minor device on another STREAMS driver. The minor device passed to the **clone** device driver during the open routine is interpreted as the major device number of another STREAMS driver for which an unused minor device is to be obtained. Each such open operation results in a separate stream to a previously unused minor device.

The **clone** device driver consists solely of an **open** subroutine. This open function performs all of the necessary work so that subsequent subroutine calls (including the **close** subroutine) require no further involvement of the **clone** device driver.

The **clone** device driver generates an **ENXIO** error, without opening the device, if the minor device number provided does not correspond to a valid major device, or if the driver indicated is not a STREAMS driver.

**Note:** Multiple opens of the same minor device cannot be done through the **clone** interface. Executing the **stat** subroutine on the file system node for a cloned device yields a different result from executing the **fstat** subroutine using a file descriptor obtained from opening the node.

## Related Information

The **close** subroutine, **fstat** subroutine, **open** subroutine, **stat** subroutine.

Understanding STREAMS Drivers and Modules and Understanding the log Device Driver in *AIX Version 6.1 Communications Programming Concepts*.

---

## copyb Utility

### Purpose

Copies a message block.

### Syntax

```
mb1k_t *  
copyb(bp)  
register mb1k_t * bp;
```

### Description

The **copyb** utility copies the contents of the message block pointed to by the *bp* parameter into a newly allocated message block of at least the same size. The **copyb** utility allocates a new block by calling the **allocb** utility. All data between the *b\_rptr* and *b\_wptr* pointers of a message block are copied to the new block, and these pointers in the new block are given the same offset values they had in the original message block.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*bp*      Contains a pointer to the message block to be copied.

### Return Values

On successful completion, the **copyb** utility returns a pointer to the new message block containing the copied data. Otherwise, it returns a null value. The copy is rounded to a fullword boundary.

## Related Information

The **allocb** utility, **copymsg** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## copymsg Utility

### Purpose

Copies a message.

## Syntax

```
mb1k_t *  
copymsg(bp)  
register mb1k_t * bp;
```

## Description

The **copymsg** utility uses the **copyb** utility to copy the message blocks contained in the message pointed to by the *bp* parameter to newly allocated message blocks. It then links the new message blocks to form the new message.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*bp*      Contains a pointer to the message to be copied.

## Return Values

On successful compilation, the **copymsg** utility returns a pointer to the new message. Otherwise, it returns a null value.

## Related Information

The **copyb** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## datamsg Utility

### Purpose

Tests whether message is a data message.

### Syntax

```
#define datamsg( type) ((type) == M_DATA | | (type) == M_PROTO | | (type) ==  
M_PCPROTO | | (type) == M_DELAY)
```

### Description

The **datamsg** utility determines if a message is a data-type message. It returns a value of True if *mp->b\_datap->db\_type* (where *mp* is declared as **mb1k\_t \*mp**) is a data-type message. The possible data types are **M\_DATA**, **M\_PROTO**, **M\_PCPROTO**, and **M\_DELAY**.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*type*      Specifies acceptable data types.

### Return Values

The **datamsg** utility returns a value of True if the message is a data-type message. Otherwise, it returns a value of False.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## dlpi STREAMS Driver

### Purpose

Provides an interface to the data link provider.

### Description

The **dlpi** driver is a STREAMS-based pseudo-driver that provides a Data Link Provider Interface (DLPI) style 2 interface to the data link providers in the operating system.

This driver is part of STREAMS Kernel Extensions.

The data link provider interface supports both the connectionless and connection-oriented modes of service, using the DL\_UNITDATA\_REQ and DL\_UNITDATA\_IND primitives. See Data Link Provider Interface Information in *AIX Version 6.1 Communications Programming Concepts*.

Refer to the "STREAMS Overview" in *AIX Version 6.1 Communications Programming Concepts* for related publications about the DLPI.

### File System Name

Each provider supported by the **dlpi** driver has a unique name in the file system. The supported interfaces are:

Driver Name	Interface
<b>/dev/dlpi/en</b>	Ethernet
<b>/dev/dlpi/et</b>	802.3
<b>/dev/dlpi/tr</b>	802.5
<b>/dev/dlpi/fi</b>	FDDI

### Physical Point of Attachment

The Physical Point of Attachment (PPA) is used to identify one of several of the same type of interface in the system. It must be a nonnegative integer in the range 0 through 99.

The **dlpi** drivers use the network interface drivers to access the communication adapter drivers. For example, the **/dev/dlpi/tr** file uses the network interface driver **if\_tr** (interface **tr0**, **tr1**, **tr2**, . . . ) to access the token-ring adapter driver. The PPA value used attaches the device open instance with the corresponding network interface. For example, opening to the **/dev/dlpi/en** device and then performing an attach with PPA value of 1 attaches this open instance to the network interface **en1**. Therefore, choosing a PPA value selects a network interface. The specific network interface must be active before a certain PPA value is used.

Examples of client and server **dlpi** programs are located in the **/usr/samples/dlpi** directory.

**Note:** You must load the **dlpi** driver using the **strload** command before running the example programs.

### Files

<b>/dev/dlpi/*</b>	Contains names of supported protocols.
<b>/usr/samples/dlpi</b>	Contains client and server <b>dlpi</b> sample programs.

## Related Information

The **ifconfig** command, **strload** command.

Understanding STREAMS Drivers and Modules, Obtaining Copies of the DLPI and TPI Specification, Data Link Provider Interface Information, in *AIX Version 6.1 Communications Programming Concepts*.

---

## dupb Utility

### Purpose

Duplicates a message-block descriptor.

### Syntax

```
mb1k_t *  
dupb(bp)  
register mb1k_t * bp;
```

### Description

The **dupb** utility duplicates the message block descriptor (**mb1k\_t**) pointed to by the *bp* parameter by copying the descriptor into a newly allocated message-block descriptor. A message block is formed with the new message-block descriptor pointing to the same data block as the original descriptor. The reference count in the data-block descriptor (**db1k\_t**) is then incremented. The **dupb** utility does not copy the data buffer, only the message-block descriptor.

Message blocks that exist on different queues can reference the same data block. In general, if the contents of a message block with a reference count greater than 1 are to be modified, the **copymsg** utility should be used to create a new message block. Only the new message block should be modified to ensure that other references to the original message block are not invalidated by unwanted changes.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*bp*     Contains a pointer to the message-block descriptor to be copied.

### Return Values

On successful compilation, the **dupb** utility returns a pointer to the new message block. If the **dupb** utility cannot allocate a new message-block descriptor, it returns a null pointer.

## Related Information

The **copymsg** utility, **dupmsg** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## dupmsg Utility

### Purpose

Duplicates a message.

## Syntax

```
mb1k_t *  
dupmsg(bp)  
register mb1k_t * bp;
```

## Description

The **dupmsg** utility calls the **dupb** utility to duplicate the message pointed to by the *bp* parameter by copying all individual message block descriptors and then linking the new message blocks to form the new message. The **dupmsg** utility does not copy data buffers, only message-block descriptors.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*bp* Specifies the message to be copied.

## Return Values

On successful completion, the **dupmsg** utility returns a pointer to the new message. Otherwise, it returns a null pointer.

## Related Information

The **dupb** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## enableok Utility

### Purpose

Enables a queue to be scheduled for service.

### Syntax

```
void  
enableok(q)  
queue_t * q;
```

### Description

The **enableok** utility cancels the effect of an earlier **noenable** utility on the same queue. It allows a queue to be scheduled for service that had previously been excluded from queue service by a call to the **noenable** utility.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue to be enabled.

### Related Information

The **noenable** utility.

## esballoc Utility

### Purpose

Allocates message and data blocks.

### Syntax

```
mb1k_t *  
esballoc(base, size, pri, free_rtn)  
unsigned char * base;  
int size, pri;  
frn_t * free_rtn;
```

### Description

The **esballoc** utility allocates message and data blocks that point directly to a client-supplied buffer. The **esballoc** utility sets the `db_base`, `b_rptr`, and `b_wptr` fields to the value specified in the *base* parameter (data buffer size) and the `db_lim` field to the *base* value plus the *size* value. The pointer to the **free\_rtn** structure is placed in the `db_freep` field of the data block.

The success of the **esballoc** utility depends on the success of the **allocb** utility and also that the *base*, *size*, and *free\_rtn* parameters are not null. If successful, the **esballoc** utility returns a pointer to a message block. If an error occurs, the **esballoc** utility returns a null pointer.

This utility is part of STREAMS Kernel Extensions.

### Parameters

<i>base</i>	Specifies the data buffer size.
<i>size</i>	Specifies the number of bytes.
<i>pri</i>	Specifies the relative importance of this block to the module. The possible values are: <ul style="list-style-type: none"><li>• <b>BPRI_LO</b></li><li>• <b>BPRI_MED</b></li><li>• <b>BPRI_HI</b></li></ul> The <i>pri</i> parameter is currently unused and is maintained only for compatibility with applications developed prior to UNIX System V Release 4.0.
<i>free_rtn</i>	Specifies the function and argument to be called when the message is freed.

### Return Values

On successful completion, the **esballoc** utility returns a pointer to a message block. Otherwise, it returns a null pointer.

### Related Information

The **allocb** utility.

---

## flushband Utility

### Purpose

Flushes the messages in a given priority band.

### Syntax

```
void flushband(q, pri, flag)
register queue_t * q;
unsigned char pri;
int flag;
```

### Description

The **flushband** utility provides modules and drivers with the capability to flush the messages associated in a given priority band. The *flag* parameter is defined the same as in the **flushq** utility. Otherwise, messages are flushed from the band specified by the *pri* parameter according to the value of the *flag* parameter.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue to flush.  
*pri* Specifies the priority band to flush. If the value of the *pri* parameter is 0, only ordinary messages are flushed.  
*flag* Specifies which messages to flush from the queue. Possible values are:

#### FLUSHDATA

Discards all **M\_DATA**, **M\_PROTO**, **M\_PCPROTO**, and **M\_DELAY** messages, but leaves all other messages on the queue.

#### FLUSHALL

Discards all messages from the queue.

### Related Information

The **flushq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## flushq Utility

### Purpose

Flushes a queue.

### Syntax

```
void flushq(q, flag)
register queue_t * q;
int flag;
```

### Description

The **flushq** utility removes messages from the message queue specified by the *q* parameter and then frees them using the **freemsg** utility.

If a queue behind the *q* parameter is blocked, the **flushq** utility may enable the blocked queue, as described in the **putq** utility.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*q* Specifies the queue to flush.  
*flag* Specifies the types of messages to flush. Possible values are:

### FLUSHDATA

Discards all **M\_DATA**, **M\_PROTO**, **M\_PCPROTO**, and **M\_DELAY** messages, but leaves all other messages on the queue.

### FLUSHALL

Discards all messages from the queue.

## Related Information

The **freemsg** utility, **putq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## freeb Utility

### Purpose

Frees a single message block.

### Syntax

```
void freeb(bp)
register struct msgb * bp;
```

### Description

The **freeb** utility frees (deallocate) the message-block descriptor pointed to by the *bp* parameter. It also frees the corresponding data block if the reference count (see the **dupb** utility) in the data-block descriptor (**datab** structure) is equal to 1. If the reference count is greater than 1, the **freeb** utility does not free the data block, but decrements the reference count instead.

If the reference count is 1 and if the message was allocated by the **esballoc** utility, the function specified by the `db_frtnp->free_func` pointer is called with the parameter specified by the `db_frtnp->free_arg` pointer.

The **freeb** utility cannot be used to free a multiple-block message (see the **freemsg** utility). Results are unpredictable if the **freeb** utility is called with a null argument. Always ensure that the pointer is nonnull before using the **freeb** utility.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*bp* Contains a pointer to the message-block descriptor that is to be freed.

## Related Information

The **dupb** utility, **esballoc** utility, **freemsg** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## freemsg Utility

### Purpose

Frees all message blocks in a message.

### Syntax

```
void freemsg(bp)
register mblk_t * bp;
```

### Description

The **freemsg** utility uses the **freeb** utility to free all message blocks and their corresponding data blocks for the message pointed to by the *bp* parameter.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*bp*      Contains a pointer to the message that is to be freed.

## Related Information

The **freeb** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## getadmin Utility

### Purpose

Returns a pointer to a module.

### Syntax

```
int
(*getadmin(mid)) ()
ushort mid;
```

### Description

The **getadmin** utility returns a pointer to the module identified by the *mid* parameter.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*mid*      Identifies the module to locate.

## Return Values

On successful completion, the **getadmin** utility returns a pointer to the specified module. Otherwise, it returns a null pointer.

## Related Information

List of Streams Programming References, Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

---

## getmid Utility

### Purpose

Returns a module ID.

### Syntax

```
ushort  
getmid(name)  
char  name;
```

### Description

The **getmid** utility returns the module ID for the module identified by the *name* parameter.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*name* Specifies the module to be identified.

## Return Values

On successful completion, the **getmid** utility returns the module ID. Otherwise, it returns a value of 0.

## Related Information

List of Streams Programming References, Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

---

## getmsg System Call

### Purpose

Gets the next message off a stream.

### Syntax

```
#include <stropts.h>  
  
int getmsg (fd, ctlptr, dataptr, flags)  
int  fd;  
struct strbuf * ctlptr;  
struct strbuf * dataptr;  
int * flags;
```

## Description

The **getmsg** system call retrieves from a STREAMS file the contents of a message located at the stream-head read queue, and places the contents into user-specified buffers. The message must contain either a data part, a control part, or both. The data and control parts of the message are placed into separate buffers, as described in the "Parameters" section. The semantics of each part are defined by the STREAMS module that generated the message.

This system call is part of the STREAMS Kernel Extensions.

## Parameters

<i>fd</i>	Specifies a file descriptor referencing an open stream.
<i>ctlptr</i>	Holds the control part of the message.
<i>dataptr</i>	Holds the data part of the message.
<i>flags</i>	Indicates the type of message to be retrieved. Acceptable values are:  <b>0</b> Process the next message of any type.  <b>RS_HIPRI</b> Process the next message only if it is a priority message.

The *ctlptr* and *dataptr* parameters each point to a **strbuf** structure that contains the following members:

```
int maxlen; /* maximum buffer length */
int len;    /* length of data */
char *buf;  /* ptr to buffer */
```

In the **strbuf** structure, the *maxlen* field indicates the maximum number of bytes this buffer can hold, the *len* field contains the number of bytes of data or control information received, and the *buf* field points to a buffer in which the data or control information is to be placed.

If the *ctlptr* (or *dataptr*) parameter is null or the *maxlen* field is -1, the following events occur:

- The control part of the message is not processed. Thus, it is left on the stream-head read queue.
- The *len* field is set to -1.

If the *maxlen* field is set to 0 and there is a zero-length control (or data) part, the following events occur:

- The zero-length part is removed from the read queue.
- The *len* field is set to 0.

If the *maxlen* field is set to 0 and there are more than 0 bytes of control (or data) information, the following events occur:

- The information is left on the read queue.
- The *len* field is set to 0.

If the *maxlen* field in the *ctlptr* or *dataptr* parameter is less than, respectively, the control or data part of the message, the following events occur:

- The *maxlen* bytes are retrieved.
- The remainder of the message is left on the stream-head read queue.
- A nonzero return value is provided.

By default, the **getmsg** system call processes the first priority or nonpriority message available on the stream-head read queue. However, a user may choose to retrieve only priority messages by setting the *flags* parameter to **RS\_HIPRI**. In this case, the **getmsg** system call processes the next message only if it

is a priority message. When the integer pointed to by *flagsp* is 0, any message will be retrieved. In this case, on return, the integer pointed to by *flagsp* will be set to **RS\_HIPRI** if a high-priority message was retrieved, or 0 otherwise.

If the **O\_NDELAY** or **O\_NONBLOCK** flag has not been set, the **getmsg** system call blocks until a message of the types specified by the *flags* parameter (priority only or either type) is available on the stream-head read queue. If the **O\_DELAY** or **O\_NONBLOCK** flag has been set and a message of the specified types is not present on the read queue, the **getmsg** system call fails and sets the **errno** global variable to **EAGAIN**.

If a hangup occurs on the stream from which messages are to be retrieved, the **getmsg** system call continues to operate until the stream-head read queue is empty. Thereafter, it returns 0 in the *len* fields of both the *ctlptr* and *dataptr* parameters.

## Return Values

Upon successful completion, the **getmsg** system call returns a nonnegative value. The possible values are:

Value	Description
0	Indicates that a full message was read successfully.
<b>MORECTL</b>	Indicates that more control information is waiting for retrieval.
<b>MOREDATA</b>	Indicates that more data is waiting for retrieval.
<b>MORECTLI.MOREDATA</b>	Indicates that both types of information remain. Subsequent <b>getmsg</b> calls retrieve the remainder of the message.

If the high priority control part of the message is consumed, the message will be placed back on the queue as a normal message of band 0. Subsequent **getmsg** system calls retrieve the remainder of the message. If, however, a priority message arrives or already exists on the **STREAM** head, the subsequent call to **getmsg** retrieves the higher-priority message before retrieving the remainder of the message that was put back.

On return, the *len* field contains one of the following:

- The number of bytes of control information or data actually received
- 0 if there is a zero-length control or data part
- -1 if no control information or data is present in the message.

If information is retrieved from a priority message, the *flags* parameter is set to **RS\_HIPRI** on return.

Upon failure, **getmsg** returns -1 and sets **errno** to indicate the error.

## Error Codes

The **getmsg** system call fails if one or more of the following is true:

Error	Description
<b>EAGAIN</b>	The <b>O_NDELAY</b> flag is set, and no messages are available.
<b>EBADF</b>	The <i>fd</i> parameter is not a valid file descriptor open for reading.
<b>EBADMSG</b>	Queued message to be read is not valid for the <b>getmsg</b> system call.
<b>EFAULT</b>	The <i>ctlptr</i> , <i>dataptr</i> , or <i>flags</i> parameter points to a location outside the allocated address space.
<b>EINTR</b>	A signal was caught during the <b>getmsg</b> system call.
<b>EINVAL</b>	An illegal value was specified in the <i>flags</i> parameter or else the stream referenced by the <i>fd</i> parameter is linked under a multiplexer.
<b>ENOSTR</b>	A stream is not associated with the <i>fd</i> parameter.

The **getmsg** system call can also fail if a STREAMS error message had been received at the stream head before the call to the **getmsg** system call. The error returned is the value contained in the STREAMS error message.

## Files

**/lib/pse.exp**                Contains the STREAMS export symbols.

## Related Information

The **poll** subroutine, **read** subroutine, **write** subroutine.

The **getpmsg** system call, **putmsg** system call, **putpmsg** system call.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getpmsg System Call

### Purpose

Gets the next priority message off a stream.

### Syntax

```
#include <stropts.h>
```

```
int getpmsg (fd, ctlptr, dataptr, bandp, flags)
int fd;
struct strbuf * ctlptr;
struct strbuf * dataptr;
int * bandp;
int * flags;
```

### Description

The **getpmsg** system call is identical to the **getmsg** system call, except that the message priority can be specified.

This system call is part of the STREAMS Kernel Extensions.

### Parameters

<i>fd</i>	Specifies a file descriptor referencing an open stream.
<i>ctlptr</i>	Holds the control part of the message.
<i>dataptr</i>	Holds the data part of the message.
<i>bandp</i>	Specifies the priority band of the message. If the value of the <i>bandp</i> parameter is set to 0, then the priority band is not limited.

*flags* Indicates the type of message priority to be retrieved. Acceptable values are:

**MSG\_ANY**

Process the next message of any type.

**MSG\_BAND**

Process the next message only if it is of the specified priority band.

**MSG\_HIPRI**

Process the next message only if it is a priority message.

If the value of the *flags* parameter is **MSG\_ANY** or **MSG\_HIPRI**, then the *bandp* parameter must be set to 0.

## Related Information

The **poll** subroutine, **read** subroutine, **write** subroutine.

The **getmsg** system call, **putmsg** system call, **putpmsg** system call.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## getq Utility

### Purpose

Gets a message from a queue.

### Syntax

```
mb1k_t *  
getq(q)  
register queue_t * q;
```

### Description

The **getq** utility gets the next available message from the queue pointed to by the *q* parameter. The **getq** utility returns a pointer to the message and removes that message from the queue. If no message is queued, the **getq** utility returns null.

The **getq** utility, and certain other utility routines, affect flow control in the Stream as follows: If the **getq** utility returns null, the queue is marked with the **QWANTR** flag so that the next time a message is placed on it, it will be scheduled for service (that is, enabled - see the **qenable** utility). If the data in the enqueued messages in the queue drops below the low-water mark, as specified by the *q\_lowat* field, and if a queue behind the current queue has previously attempted to place a message in the queue and failed, (that is, was blocked - see the **canput** utility), then the queue behind the current queue is scheduled for service.

The queue count is maintained on a per-band basis. Priority band 0 (normal messages) uses the *q\_count* and *q\_lowat* fields. Nonzero priority bands use the fields in their respective **qband** structures (the *qb\_count* and *qb\_lowat* fields). All messages appear on the same list, linked according to their *b\_next* pointers.

The *q\_count* field does not reflect the size of all messages on the queue; it only reflects those messages in the normal band of flow.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*q* Specifies the queue from which to get the message.

## Return Values

On successful completion, the **getq** utility returns a pointer to the message. Otherwise, it returns a null value.

## Related Information

The **canput** utility, **qenable** utility, **rmvq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## insq Utility

### Purpose

Puts a message at a specific place in a queue.

### Syntax

```
int
insq(q, emp, mp)
register queue_t * q;
register mblk_t * emp;
register mblk_t * mp;
```

### Description

The **insq** utility places the message pointed to by the *mp* parameter in the message queue pointed to by the *q* parameter, immediately before the already-queued message pointed to by the *emp* parameter.

If an attempt is made to insert a message out of order in a queue by using the **insq** utility, the message will not be inserted and the routine is not successful.

This utility is part of STREAMS Kernel Extensions.

The queue class of the new message is ignored. However, the priority band of the new message must adhere to the following format:

```
emp->b_prev->b_band >= mp->b_band >= emp->b_band.
```

## Parameters

*q* Specifies the queue on which to place the message.

*emp* Specifies the existing message before which the new message is to be placed.

If the *emp* parameter has a value of null, the message is placed at the end of the queue. If the *emp* parameter is nonnull, it must point to a message that exists on the queue specified by the *q* parameter, or undesirable results could occur.

*mp* Specifies the message that is to be inserted on the queue.

## Return Values

On successful completion, the **insq** utility returns a value of 1. Otherwise, it returns a value of 0.

## Related Information

The **getq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## ioctl Streams Device Driver Operations

(As defined in *X/Open Common Application Environment (CAE) Specification: System Interfaces and Headers, Issue 5 (2/97)*.)

### Purpose

Controls a STREAMS device.

### Syntax

```
#include <stropts.h>
```

```
int ioctl (fd, request, .../*arg*/)
int fd;
int request;
int .../*arg*/;
```

### Description

The **ioctl** operation performs a variety of control functions on STREAMS devices. For non-STREAMS devices, the functions performed by this call are unspecified. The *request* argument and an optional third argument (with varying type) are passed to and interpreted by the appropriate part of the STREAM associated with *fd*.

Using the **ioctl** operation on a file descriptor obtained from a call to the **shm\_open** subroutine fails with **ENOTTY**.

### Parameters

<i>fd</i>	An open file descriptor that refers to a device.
<i>request</i>	Selects the control function to be performed and will depend on the STREAMS device being addressed.
<i>.../*arg*/</i>	Represents additional information that is needed by this specific STREAMS device to perform the requested function. The type of <i>arg</i> depends on the particular control request, but it is either an integer or a pointer to a device-specific data structure.

The following **ioctl** commands, with error values indicated, are applicable to all STREAMS files:

**I\_PUSH** Pushes the module whose name is pointed to by *arg* onto the top of the current STREAM, just below the STREAM head. It then calls the *open* function of the newly-pushed module.

The *ioctl* function with the **I\_PUSH** command will fail if:

**[EINVAL]**

Invalid module name.

**[ENXIO]**

Open function of new module failed.

**[ENXIO]**

Hangup received on *fd*.

<b>I_POP</b>	<p>Removes the module just below the STREAM head of the STREAM pointed to by <i>fd</i>. The <i>arg</i> argument should be 0 in an I_POP request.</p> <p>The <i>ioctl</i> function with the I_POP command will fail if:</p> <p><b>[EINVAL]</b> No module present in the STREAM.</p> <p><b>[ENXIO]</b> Hangup received on <i>fd</i>.</p>
<b>I_LOOK</b>	<p>Retrieves the name of the module just below the STREAM head of the STREAM pointed to by <i>fd</i>, and places it in a character string pointed to by <i>arg</i>. The buffer pointed to by <i>arg</i> should be at least FMNAMESZ+1 bytes long, where FMNAMESZ is defined in <b>&lt;stropts.h&gt;</b>.</p> <p>The <i>ioctl</i> function with the I_LOOK command will fail if:</p> <p><b>[EINVAL]</b> No module present in the STREAM.</p>
<b>I_FLUSH</b>	<p>This request flushes read and/or write queues, depending on the value of <i>arg</i>. Valid <i>arg</i> values are:</p> <p><b>FLUSHR</b> Flush all read queues.</p> <p><b>FLUSHW</b> Flush all write queues.</p> <p><b>FLUSHRW</b> Flush all read and all write queues.</p> <p>The <i>ioctl</i> function with the I_FLUSH command will fail if:</p> <p><b>[EINVAL]</b> Invalid <i>arg</i> argument.</p> <p><b>[EAGAIN] or [ENOSR]</b> Unable to allocate buffers for flush messages.</p> <p><b>[ENXIO]</b> Hangup received on <i>fd</i>.</p>
<b>I_FLUSHBAND</b>	<p>Flushes a particular band of messages. The <i>arg</i> argument points to a bandinfo structure. The <b>bi_flag</b> member may be one of FLUSHR, FLUSHW, OR FLUSHRW as described above. The <b>bi_pri</b> member determines the priority band to be flushed.</p>

## I\_SETSIG

Request that the STREAMS implementation send the SIGPOLL signal to the calling process when a particular event has occurred on the STREAM associated with *fd*. I\_SETSIG supports an asynchronous processing capability in STREAMS. The value of *arg* is a bitmask that specifies the events for which the process should be signaled. It is the bitwise-OR of an combination of the following constants:

### S\_RDNORM

A normal (priority band set to 0) message has arrived at the head of a STREAM head read queue. A signal will be generated even if the message is of zero length.

### S\_RDBAND

A message with a nonzero priority band has arrived at the head of a STREAM head read queue. A signal will be generated even if the message is of zero length.

### S\_INPUT

A message, other than a high-priority message, has arrived at the head of a STREAM head read queue. A signal will be generated even if the message is of zero length.

### S\_HIPRI

A high-priority message is present on a STREAM head read queue. A signal will be generated even if the message is of zero length.

### S\_OUTPUT

The write queue for normal data (priority band 0) just below the STREAM head is no longer full. This notifies the process that there is room on the queue for sending (or writing) normal data downstream.

### S\_WRNORM

Same as S\_OUTPUT.

### S\_WRBAND

The write queue for a nonzero priority band just below the STREAM head is no longer full. This notifies the process that there is room on the queue for sending (or writing) priority data downstream.

### S\_MSG

A STREAMS signal message that contains the SIGPOLL signal has reached the front of the STREAM head read queue.

### S\_ERROR

Notification of an error condition has reached the STREAM head.

### S\_HANGUP

Notification of a hangup has reached the STREAM head.

### S\_BANDURG

When used in conjunction with S\_RDBAND, SIGURG is generated instead of SIGPOLL when a priority message reaches the front of the STREAM head read queue.

If *arg* is 0, the calling process will be unregistered and will not receive further SIGPOLL signals for the stream associated with *fd*.

Processes that wish to receive SIGPOLL signals must explicitly register to receive them using I\_SETSIG. If several processes register to receive this signal for the same event on the same STREAM, each process will be signaled when the event occurs.

The *ioctl* function with the I\_SETSIG command will fail if:

#### [EINVAL]

The value of *arg* is invalid.

#### [EINVAL]

The value of *arg* is 0 and the calling process is not registered to receive the SIGPOLL signal.

#### [EAGAIN]

There were insufficient resources to store the signal request.

<b>I_GETSIG</b>	<p>Returns the events for which the calling process is currently registered to be sent a SIGPOLL signal. The events are returned as a bitmask in an <b>int</b> pointed to by <i>arg</i>, where the events are those specified in the description of I_SETSIG above.</p> <p>The <i>ioctl</i> function with the I_GETSIG command will fail if:</p> <p><b>[EINVAL]</b> Process is not registered to receive the SIGPOLL signal.</p>
<b>I_FIND</b>	<p>The request compares the names of all modules currently present in the STREAM to the name pointed to by <i>arg</i>, and returns 1 if the named module is present in the STREAM, or returns 0 if the named module is not present.</p> <p>The <i>ioctl</i> function with the I_FIND command will fail if:</p> <p><b>[EINVAL]</b> <i>arg</i> does not contain a valid module name.</p>
<b>I_PEEK</b>	<p>This request allows a process to retrieve the information in the first message on the STREAM head read queue without taking the message off the queue. It is analogous to <i>getmsg</i> except that this command does not remove the message from the queue. The <i>arg</i> argument points to a <b>strpeek</b> structure.</p> <p>The <b>maxlen</b> member in the <b>ctlbuf</b> and <b>databuf strbuf</b> structures must be set to the number of bytes of control information and/or data information, respectively, to retrieve. The <b>flags</b> member may be marked RS_HIPRI or 0, as described by <i>getmsg</i>. If the process sets <b>flags</b> to RS_HIPRI, for example, I_PEEK will only look for a high-priority message on the STREAM head read queue.</p> <p>I_PEEK returns 1 if a message was retrieved, and returns 0 if no message was found on the STREAM head read queue, or if the RS_HIPRI flag was set in <b>flags</b> and a high-priority message was not present on the STREAM head read queue. It does not wait for a message to arrive. On return, <b>ctlbuf</b> specifies information in the control buffer, <b>databuf</b> specifies information in the data buffer, and <b>flags</b> contains the value RS_HIPRI or 0.</p>
<b>I_SRDOPT</b>	<p>Sets the read mode using the value of the argument <i>arg</i>. Read modes are described in <i>read</i>. Valid <i>arg</i> flags are:</p> <p><b>RNORM</b> Byte-stream mode, the default.</p> <p><b>RMSGD</b> Message-discard mode.</p> <p><b>RMSGN</b> Message-nondiscard mode.</p> <p>The bitwise inclusive OR of RMSGD and RMSGN will return [EINVAL]. The bitwise inclusive OR of RNORM and either RMSGD or RMSGN will result in the other flag overriding RNORM which is the default.</p> <p>In addition, treatment of control messages by the STREAM head may be changed by setting any of the following flags in <i>arg</i>:</p> <p><b>RPROTNORM</b> Fail <i>read</i> with [EBADMSG] if a message containing a control part is at the front of the STREAM head read queue.</p> <p><b>RPROTDAT</b> Deliver the control part of a message as data when a process issues a <i>read</i>.</p> <p><b>RPROTDIS</b> Discard the control part of a message, delivering any data portion, when a process issues a <i>read</i>.</p>
<b>I_GRDOPT</b>	<p>Returns the current read mode setting as, described above, in an <b>int</b> pointed to by the argument <i>arg</i>. Read modes are described in <i>read</i>.</p>

## I\_NREAD

Counts the number of data bytes in the data part of the first message on the STREAM head read queue and places this value in the **int** pointed to by *arg*. The return value for the command is the number of messages on the STREAM head read queue. For example, if 0 is returned in *arg*, but the *ioctl* return value is greater than 0, this indicates that a zero-length message is next on the queue.

## I\_FDINSERT

Creates a message from a specified buffer(s), adds information about another STREAM, and sends the message downstream. The message contains a control part and an optional data part. The data and control parts to be sent are distinguished by placement in separate buffers, as described below. The *arg* argument points to a **strfdinsert** structure.

The **len** member in the **ctlbuf strbuf** structure must be set to the size of a **t\_uscalar\_t** plus the number of bytes of control information to be sent with the message. The **fd** member specifies the file descriptor of the other STREAM, and the **offset** member, which must be suitably aligned for use as a **t\_uscalar\_t**, specifies the offset from the start of the control buffer where I\_FDINSERT will store a **t\_uscalar\_t** whose interpretation is specific to the STREAM end. The **len** member in the **databuf strbuf** structure must be set to the number of bytes of data information to be sent with the message, or to 0 if no data part is to be sent.

The **flags** member specifies the type of message to be created. A normal message is created if **flags** is set to 0, and a high-priority message is created if **flags** is set to RS\_HIPRI. For non-priority messages, I\_FDINSERT will block if the STREAM write queue is full due to internal flow control conditions. For priority messages, I\_FDINSERT does not block on this condition. For non-priority messages, I\_FDINSERT does not block when the write queue is full and O\_NONBLOCK is set. Instead, it fails and sets *errno* to [EAGAIN].

I\_FDINSERT also blocks, unless prevented by lack of internal resources, waiting for the availability of message blocks in the STREAM, regardless of priority or whether O\_NONBLOCK has been specified. No partial message is sent.

The *ioctl* function with the I\_FDINSERT command will fail if:

### [EAGAIN]

A non-priority message is specified, the O\_NONBLOCK flag is set, and the STREAM write queue is full due to internal flow control conditions.

### [EAGAIN] or [ENOSR]

Buffers cannot be allocated for the message that is to be created.

### [EINVAL]

One of the following:

- The **fd** member of the **strfdinsert** structure is not a valid, open STREAM file descriptor.
- The size of a **t\_uscalar\_t** plus *offset* is greater than the *len* member for the buffer specified through *ctlptr*.
- The *offset* member does not specify a properly-aligned location in the data buffer.
- An undefined value is stored in **flags**.

### [ENXIO]

Hangup received on the STREAM identified by either the *fd* argument or the *fd* member of the **strfdinsert** structure.

### [ERANGE]

The *len* member for the buffer specified through *databuf* does not fall within the range specified by the maximum and minimum packet sizes of the topmost STREAM module or the *len* member for the buffer specified through *databuf* is larger than the maximum configured size of the data part of a message; or the *len* member for the buffer specified through *ctlbuf* is larger than the maximum configured size of the control message.

## I\_STR

Constructs an internal STREAMS *ioctl* message from the data pointed to by *arg*, and sends that message downstream.

This mechanism is provided to send *ioctl* requests to downstream modules and drivers. It allows information to be sent with *ioctl*, and returns to the process any information sent upstream by the downstream recipient. I\_STR blocks until the system responds with either a positive or negative acknowledgment message, or until the request "times out" after some period of time. If the request times out, it fails with *errno* set to [ETIME].

At most, one I\_STR can be active on a STREAM. Further I\_STR calls will block until the active I\_STR completes at the STREAM head. The default timeout interval for these requests is 15 seconds. The O\_NONBLOCK flag has no effect on this call.

To send requests downstream, *arg* must point to a **strioc** structure.

The **ic\_cmd** member is the internal *ioctl* command intended for a downstream module or driver and **ic\_timeout** is the number of seconds:

- 1** = Infinite.
- 0** = Use implementation-dependent timeout interval.
- >0** = As specified.

an I\_STR request will wait for acknowledgment before timing out. **ic\_len** is the number of bytes in the data argument, and **ic\_dp** is a pointer to the data argument. The **ic\_len** member has two uses:

- On input, it contains the length of the data argument passed on.
- On return from the command, it contains the number of bytes being returned to the process (the buffer pointed to by **ic\_dp** should be large enough to contain the maximum amount of data that any module or the driver in the STREAM can return).

The *ioctl* function with the I\_STR command will fail if:

### [EAGAIN] or [ENOSR]

Unable to allocate buffers for the *ioctl* message.

### [EINVAL]

The *ic\_len* member is less than 0 or larger than the maximum configured size of the data part of a message, or *ic\_timeout* is less than -1.

### [ENXIO]

Hangup received on *fd*.

### [ETIME]

A downstream *ioctl* timed out before acknowledgment was received.

An I\_STR can also fail while waiting for an acknowledgment if a message indicating an error or a hangup is received at the STREAM head. In addition, an error code can be returned in the positive or negative acknowledgment message, in the event the *ioctl* command sent downstream fails. For these cases, I\_STR fails with *errno* set to the value in the message.

Sets the write mode using the value of the argument *arg*. Valid bit settings for *arg* are:

### SNDZERO

Send a zero-length message downstream when a *write* of 0 bytes occurs. To not send a zero-length message when a *write* of 0 bytes occurs, this bit must not be set in *arg* (for example, *arg* would be set to 0).

The *ioctl* function with the I\_SWROPT command will fail if:

### [EINVAL]

*arg* is not the above value.

## I\_SWROPT

## I\_GWROPT

Returns the current write mode setting, as described above, in the **int** that is pointed to by the argument *arg*.

## I\_SENDFD

I\_SENDFD creates a new reference to the open file description, associated with the file descriptor *arg*, and writes a message on the STREAMS-based pipe *fd* containing this reference, together with the user ID and group ID of the calling process.

The *ioctl* function with the I\_SENDFD command will fail if:

### [EAGAIN]

The sending STREAM is unable to allocate a message block to contain the file pointer; or the read queue of the receiving STREAM head is full and cannot accept the message sent by I\_SENDFD.

### [EBADF]

The *arg* argument is not a valid, open file descriptor.

### [EINVAL]

The *fd* argument is not connected to a STREAM pipe.

### [ENXIO]

Hangup received on *fd*.

## I\_RECVFD

Retrieves the reference to an open file description from a message written to a STREAMS-based pipe using the I\_SENDFD command, and allocates a new file descriptor in the calling process that refers to this open file description. The *arg* argument is a pointer to an **strrecvfd** data structure as defined in **stropts.h**.

The **fd** member is a file descriptor. The **uid** and **gid** members are the effective user ID and group ID, respectively, of the sending process.

If O\_NONBLOCK is not set, I\_RECVFD blocks until a message is present at the STREAM head. If O\_NONBLOCK is set, I\_RECVFD fails with *errno* set to [EAGAIN] if no message is present at the STREAM head.

If the message at the STREAM head is a message sent by an I\_SENDFD, a new file descriptor is allocated for the open file descriptor referenced in the message. The new file descriptor is placed in the **fd** member of the **strrecvfd** structure pointed to by *arg*.

The *ioctl* function with the I\_RECVFD command will fail it:

### [EAGAIN]

A message is not present at the STREAM head read queue and the O\_NONBLOCK flag is set.

### [EBADMSG]

The message at the STREAM head read queue is not a message containing a passed file descriptor.

### [EMFILE]

The process has the maximum number of file descriptors currently open that is allowed.

### [ENXIO]

Hangup received on *fd*.

## I\_LIST

This request allows the process to list all the module names on the STREAM, up to and including the topmost driver name. If *arg* is a null pointer, the return value is the number of modules, including the driver, that are on the STREAM pointed to by *fd*. This lets the process allocate enough space for the module names. Otherwise, it should point to an **str\_list** structure.

The **sl\_nmods** member indicates the number of entries that process has allocated in the array. Upon return, the **sl\_modlist** member of the **str\_list** structure contains the list of module names, and the number of entries that have been filled into the **sl\_modlist** array is found in the **sl\_nmods** member (the number includes the number of modules including the driver)> The return value from *ioctl* is 0. The entries are filled in starting at the top of the STREAM and continuing downstream until either the end of the STREAM is reached, or the number of requested modules (**sl\_nmods**) is satisfied.

The *ioctl* function with the I\_LIST command will fail if:

### [EINVAL]

The **sl\_nmods** member is less than 1.

### [EAGAIN] or [ENOSR]

Unable to allocate buffers.

## I\_ATMARK

This request allows the process to see if the message at the head of the STREAM head read queue is marked by some module downstream. The *arg* argument determines how the checking is done when there may be multiple marked messages on the STREAM head read queue. It may take on the following values:

### ANYMARK

Check if the message is marked.

### LASTMARK

Check if the message is the last one marked on the queue.

The bitwise inclusive OR of the flags ANYMARK and LASTMARK is permitted.

The return value is 1 if the mark condition is satisfied and 0 otherwise.

The *ioctl* function with the I\_ATMARK command will fail if:

### [EINVAL]

Invalid *arg* value.

## I\_CKBAND

Check if the message of given priority band exists on the STREAM head read queue. This returns 1 if a message of the given priority exists, 0 if no such message exists, or -1 on error. *arg* should be of type **int**.

The *ioctl* function with the I\_CKBAND command will fail if:

### [EINVAL]

Invalid *arg* value.

## I\_GETBAND

Return the priority band of the first message on the STREAM head read queue in the integer referenced by *arg*.

The *ioctl* function with the I\_GETBAND command will fail if:

### [ENODATA]

No message on the STREAM head read queue.

## I\_CANPUT

Check if a certain band is writable. *arg* is set to the priority band in question. The return value is 0 if the band is flow-controlled, 1 if the band is writable, or -1 on error.

The *ioctl* function with the I\_CANPUT command will fail if:

### [EINVAL]

Invalid *arg* value.

**I\_SETCLTIME** This request allows the process to set the time the STREAM head will delay when a STREAM is closing and there is no data on the write queues. Before closing each module or driver, if there is data on its write queue, the STREAM head will delay for the specified amount of time to allow the data to drain. If, after the delay, data is still present, they will be flushed. The *arg* argument is a pointer to an integer specifying the number of milliseconds to delay, rounded up to the nearest valid value. If I\_SETCLTIME is not performed on a STREAM, an implementation-dependent default timeout interval is used.

The *ioctl* function with the I\_SETCLTIME command will fail if:

**[EINVAL]**

Invalid *arg* value.

**I\_GETCLTIME** This request returns the close time delay in the integer pointed to by *arg*.

## Multiplexed STREAMS Configurations

The following four commands are used for connecting and disconnecting multiplexed STREAMS configurations. These commands use an implementation-dependent default timeout interval.

**I\_LINK** Connects two STREAMS, where *fd* is the file descriptor of the STREAM connected to the multiplexing driver, and *arg* is the file descriptor of the STREAM connected to another driver. The STREAM designated by *arg* gets connected below the multiplexing driver. I\_LINK requires the multiplexing driver to send an acknowledgment message to the STREAM head regarding the connection. This call returns a multiplexer ID number (an identifier used to disconnect the multiplexer; see I\_UNLINK) on success, and -1 on failure.

The *ioctl* function with the I\_LINK command will fail if:

**[ENXIO]**

Hangup received on *fd*.

**[ETIME]**

Time out before acknowledgment message was received at STREAM head.

**[EAGAIN] or [ENOSR]**

Unable to allocate STREAMS storage to perform the I\_LINK.

**[EINVAL]**

The *fd* argument does not support multiplexing; or *arg* is not a STREAM or is already connected downstream from a multiplexer; or the specified I\_LINK operation would connect the STREAM head in more than one place in the multiplexed STREAM.

An I\_LINK can also fail while waiting for the multiplexing driver to acknowledge the request, if a message indicating an error or a hangup is received at the STREAM head of *fd*. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, I\_LINK fails with *errno* set to the value in the message.

## I\_UNLINK

Disconnects the two STREAMS specified by *fd* and *arg*. *fd* is the file descriptor of the STREAM connected to the multiplexing driver. The *arg* argument is the multiplexer ID number that was returned by the I\_LINK *ioctl* command when a STREAM was connected downstream from the multiplexing driver. If *arg* is MUXID\_ALL, then all STREAMS that were connected to *fd* are disconnected. As in I\_LINK, this command requires acknowledgment.

The *ioctl* function with the I\_UNLINK command will fail if:

### [ENXIO]

Hangup received on *fd*.

### [ETIME]

Time out before acknowledgment message was received at STREAM head.

### [EAGAIN] or [ENOSR]

Unable to allocate buffers for the acknowledgment message.

### [EINVAL]

Invalid multiplexer ID number.

An I\_UNLINK can also fail while waiting for the multiplexing driver to acknowledge the request is a message indicating an error or a hangup is received at the STREAM head of *fd*. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, I\_UNLINK fails with *errno* set to the value in the message.

## I\_PLINK

Creates a *persistent connection* between two STREAMS, where *fd* is the file descriptor of the STREAM connected to the multiplexing driver, and *arg* is the file descriptor of the STREAM connected to another driver. This call creates a persistent connection which can exist even if the file descriptor *fd* associated with the upper STREAM to the multiplexing driver is closed. The STREAM designated by *arg* gets connected via a persistent connection below the multiplexing driver. I\_PLINK requires the multiplexing driver to send an acknowledgment to the STREAM head. This call returns a multiplexer ID number (an identifier that may be used to disconnect the multiplexer; see I\_PUNLINK) on success, and -1 on failure.

The *ioctl* function with the I\_PLINK command will fail if:

### [ENXIO]

Hangup received on *fd*.

### [ETIME]

Time out before acknowledgment message was received at STREAM head.

### [EAGAIN] or [ENOSR]

Unable to allocate STREAMS storage to perform the I\_PLINK.

### [EINVAL]

The *fd* argument does not support multiplexing; or *arg* is not a STREAM or is already connected downstream from a multiplexer; or the specified I\_PLINK operation would connect the STREAM head in more than one place in the multiplexed STREAM.

An I\_PLINK can also fail while waiting for the multiplexing driver to acknowledge the request, if a message indicating an error or a hangup is received at the STREAM head of *fd*. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, I\_PLINK fails with *errno* set to the value in the message.

**I\_PUNLINK** Disconnects the two STREAMS specified by *fd* and *arg* from a persistent connection. The *fd* argument is the file descriptor of the STREAM connected to the multiplexing driver. The *arg* argument is the multiplexer ID number that was returned by the I\_PLINK *ioctl* command when a STREAM was connected downstream from the multiplexing driver. If *arg* is MUXID\_ALL than all STREAMS which are persistent conditions to *fd* are disconnected. As in I\_PLINK, this command requires the multiplexing driver to acknowledge the request.

The *ioctl* function with the I\_PUNLINK command will fail if:

- [ENXIO]** Hangup received on *fd*.
- [ETIME]** Time out before acknowledgment message was received at STREAM head.
- [EAGAIN] or [ENOSR]** Unable to allocate buffers for the acknowledgment message.
- [EINVAL]** Invalid multiplexer ID number.

An I\_PUNLINK can also fail while waiting for the multiplexing driver to acknowledge the request if a message indicating an error or a hangup is received at the STREAM head of *fd*. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, I\_PUNLINK fails with *errno* set to the value in the message.

## Return Value

Upon successful completion, *ioctl* returns a value other than -1 that depends upon the STREAMS device control function. Otherwise, it returns -1 and sets *errno* to indicate the error.

## Errors

Under the following general conditions, *ioctl* will fail if:

- [EBADF]** The *fd* argument is not a valid open file descriptor.
- [EINTR]** A signal was caught during the *ioctl* operation.
- [EINVAL]** The STREAM or a multiplexer referenced by *fd* is linked (directly or indirectly) downstream from a multiplexer.

If an underlying device driver detects an error, then *ioctl* will fail if:

- [EINVAL]** The *request* or *arg* argument is not valid for this device.
- [EIO]** Some physical I/O error has occurred.
- [ENOTTY]** The *fd* argument is not associated with a STREAMS device that accepts control functions. A file descriptor was obtained from a call to the **shm\_open** subroutine.
- [ENXIO]** The *request* and *arg* arguments are valid for this device driver, but the service requested cannot be performed on this particular sub-device.
- [ENODEV]** The *fd* argument refers to a valid STREAMS device, but the corresponding device driver does not support the *ioctl* function.

If a STREAM is connected downstream from a multiplexer, and *ioctl* command except I\_UNLINK and I\_PUNLINK will set *errno* to [EINVAL].

## Application Usage

The implementation-dependent timeout interval for STREAMS has historically been 15 seconds.

## Related Information

The **close** subroutine, **getmsg** system call, **open** subroutine, **poll** subroutine, **putmsg** system call, **read** subroutine, **write** subroutine

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## I\_ATMARK streamio Operation

### Purpose

Checks to see if a message is marked.

### Description

The **I\_ATMARK** operation shows the user if the current message on the stream-head read queue is marked by a downstream module. The *arg* parameter determines how the checking is done when there are multiple marked messages on the stream-head read queue. The possible values for the *arg* parameter are:

Value	Description
<b>ANYMARK</b>	Read to determine if the message is marked by a downstream module.
<b>LASTMARK</b>	Read to determine if the message is the last one marked on the queue by a downstream module.

The **I\_ATMARK** operation returns a value of 1 if the mark condition is satisfied. Otherwise, it returns a value of 0.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EINVAL</b>	The value of the <i>arg</i> parameter could not be used.

---

## I\_CANPUT streamio Operation

### Purpose

Checks if a given band is writable.

### Description

The **I\_CANPUT** operation checks a given priority band to see if it can be written on. The *arg* parameter contains the priority band to be checked.

### Return Values

The return value is set to one of the following:

Value	Description
<b>0</b>	The band is flow controlled.
<b>1</b>	The band is writable.
<b>-1</b>	An error occurred.

## Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EINVAL</b>	The value in the <i>arg</i> parameter is invalid.

---

## I\_CKBAND streamio Operation

### Purpose

Checks if a message of a particular band is on the stream-head read queue.

### Description

The **I\_CKBAND** operation checks to see if a message of a given priority band exists on the stream-head read queue. The *arg* parameter is an integer containing the value of the priority band being searched for.

The **I\_CKBAND** operation returns a value of 1 if a message of the given band exists. Otherwise, it returns a value of -1.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EINVAL</b>	The value in the <i>arg</i> parameter is not valid.

---

## I\_FDINSERT streamio Operation

### Purpose

Creates a message from user-specified buffers, adds information about another stream and sends the message downstream.

### Description

The **I\_FDINSERT** operation creates a message from user-specified buffers, adds information about another stream, and sends the message downstream. The message contains a control part and an optional data part. The data and control parts transmitted are identified by their placement in separate buffers. The *arg* parameter points to a **strfdinsert** structure that contains the following elements:

```
struct strbuf  ctlbuf;
struct strbuf  databuf;
long          flags;
int           fildes;
int           offset;
```

The *len* field in the **strbuf** structure must be set to the size of a pointer plus the number of bytes of control information sent with the message. The *fildes* field in the **strfdinsert** structure specifies the file descriptor of the other stream. The *offset* field, which must be word-aligned, specifies the number of bytes beyond the beginning of the control buffer to store a pointer. This pointer will be the address of the read queue structure of the driver for the stream corresponding to the *fildes* field in the **strfdinsert** structure. The *len* field in the **strbuf** structure of the *databuf* field must be set to the number of bytes of data information sent with the message or to 0 if no data part is sent.

The `flags` field specifies the type of message created. There are two valid values for the `flags` field:

Value	Description
0	Creates a nonpriority message.
RS_HIPRI	Creates a priority message.

For nonpriority messages, the **I\_FDINSERT** operation blocks if the stream write queue is full due to internal flow-control conditions. For priority messages, the **I\_FDINSERT** operation does not block on this condition. For nonpriority messages, the **I\_FDINSERT** operation does not block when the write queue is full and the **O\_NDELAY** flag is set. Instead, the operation fails and sets the **errno** global variable to **EAGAIN**.

The **I\_FDINSERT** operation also blocks unless prevented by lack of internal resources, while it is waiting for the availability of message blocks in the stream, regardless of priority or whether the **O\_NDELAY** flag has been specified. No partial message is sent.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EAGAIN</b>	A nonpriority message was specified, the <b>O_NDELAY</b> flag is set, and the stream write queue is full due to internal flow-control conditions.
<b>ENOSR</b>	Buffers could not be allocated for the message that was to be created due to insufficient STREAMS memory resources.
<b>EFAULT</b>	The <i>arg</i> parameter points to an area outside the allocated address space, or the buffer area specified in the <code>ctlbuf</code> or <code>databuf</code> field is outside this space.
<b>EINVAL</b>	One of the following conditions has occurred:  The <code>fildev</code> field in the <b>strfdinsert</b> structure is not a valid, open stream file descriptor.  The size of a pointer plus the value of the <code>offset</code> field is greater than the <code>len</code> field for the buffer specified through the <code>ctlptr</code> field.  The <code>offset</code> parameter does not specify a properly aligned location in the data buffer.  An undefined value is stored in the <code>flags</code> parameter.
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter of the <b>ioctl</b> call or the <code>fildev</code> field in the <b>strfdinsert</b> structure.
<b>ERANGE</b>	The <code>len</code> field for the buffer specified through the <code>databuf</code> field does not fall within the range specified by the maximum and minimum packet sizes of the topmost stream module; or the <code>len</code> field for the buffer specified through the <code>databuf</code> field is larger than the maximum configured size of the data part of a message; or the <code>len</code> field for the buffer specified through the <code>ctlbuf</code> field is larger than the maximum configured size of the control part of a message.

The **I\_FDINSERT** operation is also unsuccessful if an error message is received by the stream head corresponding to the `fildev` field in the **strfdinsert** structure. In this case, the **errno** global variable is set to the value in the message.

---

## I\_FIND streamio Operation

### Purpose

Compares the names of all modules currently present in the stream to a specified name.

## Description

The **I\_FIND** operation compares the names of all modules currently present in the stream to the name pointed to by the *arg* parameter, and returns a value of 1 if the named module is present in the stream. It returns a value of 0 if the named module is not present.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.
<b>EINVAL</b>	The <i>arg</i> parameter does not contain a valid module name.

---

## I\_FLUSH streamio Operation

### Purpose

Flushes all input or output queues.

### Description

The **I\_FLUSH** operation flushes all input or output queues, depending on the value of the *arg* parameter. Legal values for the *arg* parameter are:

Value	Description
<b>FLUSHR</b>	Flush read queues.
<b>FLUSHW</b>	Flush write queues.
<b>FLUSHRW</b>	Flush read and write queues.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENOSR</b>	Unable to allocate buffers for flush message due to insufficient STREAMS memory resources.
<b>EINVAL</b>	Invalid value for the <i>arg</i> parameter.
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter.

---

## I\_FLUSHBAND streamio Operation

### Purpose

Flushes all messages from a particular band.

### Description

The **I\_FLUSHBAND** operation flushes all messages of a given priority band from all input or output queues. The *arg* parameter points to a **bandinfo** structure that contains the following elements:

```
unsigned char  bi_pri;  
int           bi_flag;
```

The elements are defined as follows:

Element	Description
<code>bi_pri</code>	Specifies the band to be flushed.
<code>bi_flag</code>	Specifies the queues to be pushed. Legal values for the <code>bi_flag</code> field are:  <b>FLUSHR</b> Flush read queues.  <b>FLUSHW</b> Flush write queues.  <b>FLUSHRW</b> Flush read and write queues.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENOSR</b>	Unable to allocate buffers for flush message due to insufficient STREAMS memory resources.
<b>EINVAL</b>	Invalid value for the <i>arg</i> parameter.
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter.

---

## I\_GETBAND streamio Operation

### Purpose

Gets the band of the first message on the stream-head read queue.

### Description

The **I\_GETBAND** operation returns the priority band of the first message on the stream-head read queue in the integer referenced by the *arg* parameter.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>ENODATA</b>	No message is on the stream-head read queue.

---

## I\_GETCLTIME streamio Operation

### Purpose

Returns the delay time.

### Description

The **I\_GETCLTIME** operation returns the delay time, in milliseconds, that is pointed to by the *arg* parameter.

This operation is part of STREAMS Kernel Extensions.

---

## I\_GETSIG streamio Operation

### Purpose

Returns the events for which the calling process is currently registered to be sent a **SIGPOLL** signal.

### Description

The **I\_GETSIG** operation returns the events for which the calling process is currently registered to be sent a **SIGPOLL** signal. The events are returned as a bitmask pointed to by the *arg* parameter, where the events are those specified in the description of the **I\_SETSIG** operation.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EINVAL</b>	Process not registered to receive the <b>SIGPOLL</b> signal.
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.

---

## I\_GRDOPT streamio Operation

### Purpose

Returns the current read mode setting.

### Description

The **I\_GRDOPT** operation returns the current read mode setting in an *int* parameter pointed to by the *arg* parameter. Read modes are described in the **read** subroutine description.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.

---

## I\_LINK streamio Operation

### Purpose

Connects two specified streams.

### Description

The **I\_LINK** operation is used for connecting multiplexed STREAMS configurations.

The **I\_LINK** operation connects two streams, where the *fildev* parameter is the file descriptor of the stream connected to the multiplexing driver, and the *arg* parameter is the file descriptor of the stream connected to another driver. The stream designated by the *arg* parameter gets connected below the multiplexing

driver. The **I\_LINK** operation requires the multiplexing driver to send an acknowledgment message to the stream head regarding the linking operation. This call returns a multiplexer ID number (an identifier used to disconnect the multiplexer; see the **I\_UNLINK** operation) on success, and a value of -1 on failure.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENXIO</b>	Hangup received on the <i>fildev</i> field.
<b>ETIME</b>	Time out before acknowledgment message was received at stream head.
<b>EAGAIN</b>	Temporarily unable to allocate storage to perform the <b>I_LINK</b> operation.
<b>ENOSR</b>	Unable to allocate storage to perform the <b>I_LINK</b> operation due to insufficient STREAMS memory resources.
<b>EBADF</b>	The <i>arg</i> parameter is not a valid, open file descriptor.
<b>EINVAL</b>	The specified link operation would cause a cycle in the resulting configuration; that is, if a given stream head is linked into a multiplexing configuration in more than one place.

An **I\_LINK** operation can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the stream head of the *fildev* parameter. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, the **I\_LINK** operation fails with the **errno** global variable set to the value in the message.

---

## I\_LIST streamio Operation

### Purpose

Lists all the module names on a stream.

### Description

The **I\_LIST** operation lists all of the modules present on a stream, including the topmost driver name. If the value of the *arg* parameter is null, the **I\_LIST** operation returns the number of modules on the stream pointed to by the *fildev* parameter. If the value of the *arg* parameter is nonnull, it points to an **str\_list** structure that contains the following elements:

```
int sl_nmods;
struct str_mlist *sl_modlist;
```

The **str\_mlist** structure contains the following element:

```
char l_name[FMNAMESZ+1];
```

The fields are defined as follows:

Field	Description
<i>sl_nmods</i>	Specifies the number of entries the user has allocated in the array.
<i>sl_modlist</i>	Contains the list of module names (on return).

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EAGAIN</b>	Unable to allocate buffers.
<b>EINVAL</b>	The <code>sl_nmods</code> member is less than 1.

---

## I\_LOOK streamio Operation

### Purpose

Retrieves the name of the module just below the stream head.

### Syntax

```
#include <sys/conf.h>
#include <stropts.h>

int ioctl (fildes, command, arg)
int fildes, command;
```

### Description

The **I\_LOOK** operation retrieves the name of the module just below the stream head of the stream pointed to by the *fildes* parameter and places it in a null terminated character string pointed at by the *arg* parameter. The buffer pointed to by the *arg* parameter should be at least `FMNAMESMZ + 1` bytes long.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.
<b>EINVAL</b>	No module is present in stream.

---

## I\_NREAD streamio Operation

### Purpose

Counts the number of data bytes in data blocks in the first message on the stream-head read queue, and places this value in a specified location.

### Description

The **I\_NREAD** operation counts the number of data bytes in data blocks in the first message on the stream-head read queue, and places this value in the location pointed to by the *arg* parameter.

This operation is part of STREAMS Kernel Extensions.

### Return Values

The return value for the operation is the number of messages on the stream-head read queue. For example, if a value of 0 is returned in the *arg* parameter, but the **ioctl** operation return value is greater than 0, this indicates that a zero-length message is next on the queue.

## Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.

---

## I\_PEEK streamio Operation

### Purpose

Allows a user to retrieve the information in the first message on the stream-head read queue without taking the message off the queue.

### Description

The **I\_PEEK** operation allows a user to retrieve the information in the first message on the stream-head read queue without taking the message off the queue. The *arg* parameter points to a **strpeek** structure that contains the following elements:

```
struct strbuf ctlbuf;  
struct strbuf databuf;  
long flags;
```

The *maxlen* field in the **strbuf** structures of the *ctlbuf* and *databuf* fields must be set to the number of bytes of control information or data information, respectively, to retrieve. If the user sets the *flags* field to **RS\_HIPRI**, the **I\_PEEK** operation looks for a priority message only on the stream-head read queue.

The **I\_PEEK** operation returns a value of 1 if a message was retrieved, and returns a value of 0 if no message was found on the stream-head read queue, or if the **RS\_HIPRI** flag was set in the *flags* field and a priority message was not present on the stream-head read queue. It does not wait for a message to arrive.

On return, the fields contain the following data:

Data	Description
<i>ctlbuf</i>	Specifies information in the control buffer.
<i>databuf</i>	Specifies information in the data buffer.
<i>flags</i>	Contains the value of 0 or <b>RS_HIPRI</b> .

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points, or the buffer area specified in the <i>ctlbuf</i> or <i>databuf</i> field is outside the allocated address space.
<b>EBADMSG</b>	Queued message is not valid for the <b>I_PEEK</b> operation.

---

## I\_PLINK streamio Operation

### Purpose

Connects two specified streams.

## Description

The **I\_PLINK** operation is used for connecting multiplexed STREAMS configurations with a permanent link.

This operation is part of STREAMS Kernel Extensions.

The **I\_PLINK** operation connects two streams, where the *fildev* parameter is the file descriptor of the stream connected to the multiplexing driver, and the *arg* parameter is the file descriptor of the stream connected to another driver. The stream designated by the *arg* parameter gets connected by a permanent link below the multiplexing driver. The **I\_PLINK** operation requires the multiplexing driver to send an acknowledgment message to the stream head regarding the linking operation. This call creates a permanent link which can exist even if the file descriptor associated with the upper stream to the multiplexing driver is closed. This call returns a multiplexer ID number (an identifier used to disconnect the multiplexer; see the **I\_PUNLINK** operation) on success, and a value of -1 on failure.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENXIO</b>	Hangup received on the <i>fildev</i> field.
<b>ETIME</b>	Time out occurred before acknowledgment message was received at stream head.
<b>EAGAIN</b>	Unable to allocate storage to perform the <b>I_PLINK</b> operation.
<b>EBADF</b>	The <i>arg</i> parameter is not a valid, open file descriptor.
<b>EINVAL</b>	The <i>fildev</i> parameter does not support multiplexing.
OR	
The <i>fildev</i> parameter is the file descriptor of a pipe or FIFO.	
OR	
The <i>arg</i> parameter is not a stream or is already linked under a multiplexer.	
OR	
The specified link operation would cause a cycle in the resulting configuration; that is, if a given stream head is linked into a multiplexing configuration in more than one place.	

An **I\_PLINK** operation can also be unsuccessful while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the stream head of the *fildev* parameter. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, the **I\_PLINK** operation is unsuccessful with the **errno** global variable set to the value in the message.

---

## I\_POP streamio Operation

### Purpose

Removes the module just below the stream head.

### Description

The **I\_POP** operation removes the module just below the stream head of the stream pointed to by the *fildev* parameter. The value of the *arg* parameter should be 0 in an **I\_POP** request.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EINVAL</b>	No module is present in the stream.
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter.

---

## I\_PUNLINK streamio Operation

### Purpose

Disconnects the two specified streams.

### Description

The **I\_PUNLINK** operation is used for disconnecting Multiplexed STREAMS configurations connected by a permanent link.

The **I\_PUNLINK** operation disconnects the two streams specified by the *fildev* parameter and the *arg* parameter that are connected with a permanent link. The *fildev* parameter is the file descriptor of the stream connected to the multiplexing driver. The *arg* parameter is the multiplexer ID number that was returned by the **I\_PLINK** operation. If the value of the *arg* parameter is **MUXID\_ALL**, then all streams which are permanently linked to the stream specified by the *fildev* parameter are disconnected. As in the **I\_PLINK** operation, this operation requires the multiplexing driver to acknowledge the unlink.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter.
<b>ETIME</b>	Time out occurred before acknowledgment message was received at stream head.
<b>EINVAL</b>	The <i>arg</i> parameter is an invalid multiplexer ID number.

OR

The *fildev* parameter is the file descriptor of a pipe or FIFO.

An **I\_PUNLINK** operation can also be unsuccessful while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the stream head of the *fildev* parameter. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, the **I\_PUNLINK** operation is unsuccessful and the **errno** global variable is set to the value in the message.

---

## I\_PUSH streamio Operation

### Purpose

Pushes a module onto the top of the current stream.

### Description

The **I\_PUSH** operation pushes the module whose name is pointed to by the *arg* parameter onto the top of the current stream, just below the stream head. It then calls the open routine of the newly-pushed module.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EINVAL</b>	Incorrect module name.
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.
<b>ENXIO</b>	Open routine of new module failed.
<b>ENXIO</b>	Hangup received on the <i>fildev</i> parameter.

---

## I\_RECVFD streamio Operation

### Purpose

Retrieves the file descriptor associated with the message sent by an **I\_SENDFD** operation over a stream pipe.

### Description

The **I\_RECVFD** operation retrieves the file descriptor associated with the message sent by an **I\_SENDFD** operation over a stream pipe. The *arg* parameter is a pointer to a data buffer large enough to hold an **strrecvfd** data structure containing the following elements:

```
int fd;
unsigned short uid;
unsigned short gid;
char fill[8];
```

The fields of the **strrecvfd** structure are defined as follows:

Field	Description
<i>fd</i>	Specifies an integer file descriptor.
<i>uid</i>	Specifies the user ID of the sending stream.
<i>gid</i>	Specifies the group ID of the sending stream.

If the **O\_NDELAY** flag is not set, the **I\_RECVFD** operation blocks until a message is present at the stream head. If the **O\_NDELAY** flag is set, the **I\_RECVFD** operation fails with the **errno** global variable set to **EAGAIN** if no message is present at the stream head.

If the message at the stream head is a message sent by an **I\_SENDFD** operation, a new user file descriptor is allocated for the file pointer contained in the message. The new file descriptor is placed in the *fd* field of the **strrecvfd** structure. The structure is copied into the user data buffer pointed to by the *arg* parameter.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EAGAIN</b>	A message was not present at the stream head read queue, and the <b>O_NDELAY</b> flag is set.
<b>EBADMSG</b>	The message at the stream head read queue was not a message containing a passed file descriptor.
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.
<b>EMFILE</b>	The <b>NOFILES</b> file descriptor is currently open.

Value	Description
ENXIO	Hangup received on the <i>fildev</i> parameter.

---

## I\_SENDFD streamio Operation

### Purpose

Requests a stream to send a message to the stream head at the other end of a stream pipe.

### Description

The **I\_SENDFD** operation requests the stream associated with the *fildev* field to send a message, containing a file pointer, to the stream head at the other end of a stream pipe. The file pointer corresponds to the *arg* parameter, which must be an integer file descriptor.

The **I\_SENDFD** operation converts the *arg* parameter into the corresponding system file pointer. It allocates a message block and inserts the file pointer in the block. The user ID and group ID associated with the sending process are also inserted. This message is placed directly on the read queue of the stream head at the other end of the stream pipe to which it is connected.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
EAGAIN	The sending stream is unable to allocate a message block to contain the file pointer.
EAGAIN	The read queue of the receiving stream head is full and cannot accept the message sent by the <b>I_SENDFD</b> operation.
EBADF	The <i>arg</i> parameter is not a valid, open file descriptor.
EINVAL	The <i>fildev</i> parameter is not connected to a stream pipe.
ENXIO	Hangup received on the <i>fildev</i> parameter.

---

## I\_SETCLTIME streamio Operation

### Purpose

Sets the time that the stream head delays when a stream is closing.

### Description

The **I\_SETCLTIME** operation sets the time that the stream head delays when a stream is closing and there is data on the write queues. Before closing each module and driver, the stream head delays closing for the specified length of time to allow the data to be written. Any data left after the delay is flushed.

The *arg* parameter contains a pointer to the number of milliseconds to delay. This number is rounded up to the nearest legal value on the system. The default delay time is 15 seconds.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EINVAL</b>	The value in the <i>arg</i> parameter is invalid.

---

## I\_SETSIG streamio Operation

### Purpose

Informs the stream head that the user wishes the kernel to issue the **SIGPOLL** signal when a particular event occurs on the stream.

### Description

The **I\_SETSIG** operation informs the stream head that the user wishes the kernel to issue the **SIGPOLL** signal (see the **signal** and **sigset** subroutines) when a particular event has occurred on the stream associated with the *fildev* parameter. The **I\_SETSIG** operation supports an asynchronous processing capability in STREAMS. The value of the *arg* parameter is a bit mask that specifies the events for which the user should be signaled. It is the bitwise-OR of any combination of the following constants:

Constant	Description
<b>S_INPUT</b>	A nonpriority message has arrived on a stream-head read queue, and no other messages existed on that queue before this message was placed there. This is set even if the message is of zero length.
<b>S_HIPRI</b>	A priority message is present on the stream-head read queue. This is set even if the message is of zero length.
<b>S_OUTPUT</b>	The write queue just below the stream head is no longer full. This notifies the user that there is room on the queue for sending (or writing) data downstream.
<b>S_MSG</b>	A STREAMS signal message that contains the <b>SIGPOLL</b> signal has reached the front of the stream-head read queue.

A user process may choose to be signaled only by priority messages by setting the *arg* bit mask to the value **S\_HIRPI**.

Processes that wish to receive **SIGPOLL** signals must explicitly register to receive them using **I\_SETSIG**. If several processes register to receive this signal for the same event on the same stream, each process will be signaled when the event occurs.

If the value of the *arg* parameter is 0, the calling process is unregistered and does not receive further **SIGPOLL** signals.

This operation is part of STREAMS Kernel Extensions.

### Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EINVAL</b>	The value for the <i>arg</i> parameter is invalid or 0 and process is not registered to receive the <b>SIGPOLL</b> signal.
<b>EAGAIN</b>	The allocation of a data structure to store the signal request is unsuccessful.

---

## I\_SRDOPT streamio Operation

### Purpose

Sets the read mode.

### Description

The **I\_SRDOPT** operation sets the read mode using the value of the *arg* parameter. Legal values for the *arg* parameter are:

Value	Description
<b>RNORM</b>	Byte-stream mode. This is the default mode.
<b>RMSGD</b>	Message-discard mode.
<b>RMSGN</b>	Message-nondiscard mode.
<b>RFILL</b>	Read mode. This mode prevents completion of any <b>read</b> request until one of three conditions occurs: <ul style="list-style-type: none"><li>• The entire user buffer is filled.</li><li>• An end of file occurs.</li><li>• The stream head receives an <b>M_MI_READ_END</b> message.</li></ul>

Several control messages support the **RFILL** mode. They are used by modules to manipulate data being placed in user buffers at the stream head. These messages are multiplexed under a single **M\_MI** message type. The message subtype, pointed to by the *b\_rptr* parameter, is one of the following:

#### **M\_MI\_READ\_SEEK**

Provides random access data retrieval. An application and a cooperating module can gather large data blocks from a slow, high-latency, or unreliable link, while minimizing the number of system calls required, and relieving the protocol modules of large buffering requirements.

The **M\_MI\_READ\_SEEK** message subtype is followed by two long words, as in a standard **seek** call. The first word is an origin indicator as follows:

<b>0</b>	Start of buffer
<b>1</b>	Current position
<b>2</b>	End of buffer

The second word is a signed offset from the specified origin.

#### **M\_MI\_READ\_RESET**

Discards any data previously delivered to partially satisfy an **RFILL** mode **read** request.

#### **M\_MI\_READ\_END**

Completes the current **RFILL** mode **read** request with whatever data has already been delivered.

In addition, treatment of control messages by the stream head can be changed by setting the following flags in the *arg* parameter:

Flag	Description
<b>RPROTNORM</b>	Causes the <b>read</b> routine to be unsuccessful if a control message is at the front of the stream-head read queue.
<b>RPROTDAT</b>	Delivers the control portion of a message as data.
<b>RPROTDIS</b>	Discards the control portion of a message, delivering any data portion.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
<b>EINVAL</b>	The value of the <i>arg</i> parameter is not one of the above legal values.

---

## I\_STR streamio Operation

### Purpose

Constructs an internal STREAMS ioctl message.

### Description

The **I\_STR** operation constructs an internal STREAMS ioctl message from the data pointed to by the *arg* parameter and sends that message downstream.

This mechanism is provided to send user ioctl requests to downstream modules and drivers. It allows information to be sent with the ioctl and returns to the user any information sent upstream by the downstream recipient. The **I\_STR** operation blocks until the system responds with either a positive or negative acknowledgment message or until the request times out after some period of time. If the request times out, it fails with the **errno** global variable set to **ETIME**.

At most, one **I\_STR** operation can be active on a stream. Further **I\_STR** operation calls block until the active **I\_STR** operation completes at the stream head. The default timeout interval for this request is 15 seconds. The **O\_NDELAY** flag has no effect on this call.

To send a request downstream, the *arg* parameter must point to a **strioc** structure that contains the following elements:

```
int ic_cmd;      /* downstream operation */
int ic_timeout;  /* ACK/NAK timeout */
int ic_len;      /* length of data arg */
char *ic_dp;     /* ptr to data arg */
```

The elements of the **strioc** structure are described as follows:

Element	Description
ic_cmd	The internal <b>ioctl</b> operation intended for a downstream module or driver.
ic_timeout	The number of seconds an <b>I_STR</b> request waits for acknowledgment before timing out: <ul style="list-style-type: none"><li><b>-1</b> Waits an infinite number of seconds.</li><li><b>0</b> Uses default value.</li><li><b>&gt; 0</b> Waits the specified number of seconds.</li></ul>
ic_len	The number of bytes in the data argument. The <i>ic_len</i> field has two uses: <ul style="list-style-type: none"><li>On input, it contains the length of the data argument passed in.</li><li>On return from the operation, it contains the number of bytes being returned to the user (the buffer pointed to by the <i>ic_dp</i> field should be large enough to contain the maximum amount of data that any module or the driver in the stream can return).</li></ul>
ic_dp	A pointer to the data parameter.

The stream head converts the information pointed to by the **strioc** structure to an internal **ioctl** operation message and sends it downstream.

This operation is part of STREAMS Kernel Extensions.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>EAGAIN</b>	The value of <b>ic_len</b> is greater than the maximum size of a message block returned by the STREAMS <b>allocb</b> utility, or there is insufficient memory for a message block.
<b>ENOSR</b>	Unable to allocate buffers for the <b>ioctl</b> message due to insufficient STREAMS memory resources.
<b>EFAULT</b>	The area pointed to by the <b>arg</b> parameter or the buffer area specified by the <b>ic_dp</b> and <b>ic_len</b> fields (for data sent and data returned, respectively) is outside of the allocated address space.
<b>EINVAL</b>	The value of the <b>ic_len</b> field is less than 0 or greater than the maximum configured size of the data part of a message, or the value of the <b>ic_timeout</b> field is less than -1.
<b>ENXIO</b>	Hangup received on the <b>fildev</b> field.
<b>ETIME</b>	A downstream <b>streamio</b> operation timed out before acknowledgment was received.

An **I\_STR** operation can also be unsuccessful while waiting for an acknowledgment if a message indicating an error or a hangup is received at the stream head. In addition, an error code can be returned in the positive or negative acknowledgment messages, in the event that the **streamio** operation sent downstream fails. For these cases, the **I\_STR** operation is unsuccessful and the **errno** global variable is set to the value in the message.

---

## I\_UNLINK streamio Operation

### Purpose

Disconnects the two specified streams.

### Description

The **I\_UNLINK** operation is used for disconnecting multiplexed STREAMS configurations.

This operation is part of STREAMS Kernel Extensions.

The **I\_UNLINK** operation disconnects the two streams specified by the **fildev** parameter and the **arg** parameter. The **fildev** parameter is the file descriptor of the stream connected to the multiplexing driver. The **fildev** parameter must correspond to the stream on which the **ioctl I\_LINK** operation was issued to link the stream below the multiplexing driver. The **arg** parameter is the multiplexer ID number that was returned by the **I\_LINK** operation. If the value of the **arg** parameter is -1, then all streams that were linked to the **fildev** parameter are disconnected. As in the **I\_LINK** operation, this operation requires the multiplexing driver to acknowledge the unlink.

## Error Codes

If unsuccessful, the **errno** global variable is set to one of the following values:

Value	Description
<b>ENXIO</b>	Hangup received on the <b>fildev</b> parameter.
<b>ETIME</b>	Time out before acknowledgment message was received at stream head.
<b>ENOSR</b>	Unable to allocate storage to perform the <b>I_UNLINK</b> operation due to insufficient STREAMS memory resources.
<b>EINVAL</b>	The <b>arg</b> parameter is an invalid multiplexer ID number or the <b>fildev</b> parameter is not the stream on which the <b>I_LINK</b> operation that returned the <b>arg</b> parameter was performed.

An **I\_UNLINK** operation can also fail while waiting for the multiplexing driver to acknowledge the link request, if a message indicating an error or a hangup is received at the stream head of the **fildev**

parameter. In addition, an error code can be returned in the positive or negative acknowledgment message. For these cases, the **I\_UNLINK** operation fails and the **errno** global variable is set to the value in the message.

## Related Information

The **I\_LINK** streamio operation, **I\_PUNLINK** streamio operation.

List of Streams Programming References, Understanding streamio (STREAMS ioctl) Operations, Understanding STREAMS Drivers and Modules, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## isastream Function

### Purpose

Tests a file descriptor.

### Library

Standard C Library (**libc.a**)

### Syntax

```
int isastream(int fildev);
```

### Description

The **isastream** subroutine determines if a file descriptor represents a STREAMS file.

### Parameters

*fildev*            Specifies which open file to check.

### Return Values

On successful completion, the **isastream** subroutine returns a value of 1 if the *fildev* parameter represents a STREAMS file, or a value of 0 if not. Otherwise, it returns a value of -1 and sets the **errno** global variable to indicate the error.

### Error Codes

If unsuccessful, the **errno** global variable is set to the following value:

Value	Description
EBADF	The <i>fildev</i> parameter does not specify a valid open file.

## Related Information

**streamio** operations.

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts*.

---

## linkb Utility

### Purpose

Concatenates two messages into one.

## Syntax

```
void link(mp, bp)  
register mblk_t * mp;  
register mblk_t * bp;
```

## Description

The **linkb** utility puts the message pointed to by the *bp* parameter at the tail of the message pointed to by the *mp* parameter. This results in a single message.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*mp* Specifies the message to which the second message is to be linked.  
*bp* Specifies the message that is to be linked to the end of first message.

## Related Information

The **unlinkb** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## mi\_bufcall Utility

### Purpose

Provides a reliable alternative to the **bufcall** utility.

### Syntax

```
#include <pse/mi.h>  
#include <sys/stream.h>  
  
void mi_bufcall ( Queue, Size, Priority)  
queue_t *Queue;  
int Size;  
int Priority;
```

### Description

The **mi\_bufcall** utility provides a reliable alternative to the **bufcall** utility. The standard STREAMS **bufcall** utility is intended to be called when the **allocb** utility is unable to allocate a block for a message, and invokes a specified callback function (typically the **qenable** utility) with a given queue when a large enough block becomes available. This can cause system problems if the stream closes so that the queue becomes invalid before the callback function is invoked.

The **mi\_bufcall** utility is a reliable alternative, as the queue is not deallocated until the call is complete. This utility uses the standard **bufcall** mechanism with its own internal callback routine. The callback routine either invokes the **qenable** utility with the specified *Queue* parameter, or simply deallocates the instance data associated with the stream if the queue has already been closed.

The **mi\_bufcall** utility is part of STREAMS kernel extensions.

**Note:** The **stream.h** header file must be the last included header file of each source file using the stream library.

## Parameters

<i>Queue</i>	Specifies the queue which is to be passed to the <b>qenable</b> utility.
<i>Size</i>	Specifies the required buffer size.
<i>Priority</i>	Specifies the priority as used by the standard STREAMS <b>bufcall</b> mechanism.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts* .

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts* .

The **bufcall** utility, **mi\_close\_comm** utility, **mi\_next\_ptr** utility, **mi\_open\_comm** utility.

---

## mi\_close\_comm Utility

### Purpose

Performs housekeeping during STREAMS driver or module close operations.

### Syntax

```
#include <pse/mi.h>
#include <sys/stream.h>

int mi_close_comm ( StaticPointer, Queue)
caddr_t *StaticPointer;
queue_t *Queue;
```

### Description

The **mi\_close\_comm** utility performs housekeeping during STREAMS driver or module close operations. It is intended to be called by the driver or module **close** routine. It releases the memory allocated by the corresponding call to the **mi\_open\_comm** utility, and frees the minor number for reuse.

If an **mi\_bufcall** operation is outstanding, module resources are not freed until the **mi\_bufcall** operation is complete.

The **mi\_close\_comm** utility is part of STREAMS kernel extensions.

#### Notes:

1. Each call to the **mi\_close\_comm** utility must have a corresponding call to the **mi\_open\_comm** utility. Executing one of these utilities without making a corresponding call to the other will lead to unpredictable results.
2. The **stream.h** header file must be the last included header file of each source file using the stream library.

## Parameters

<i>StaticPointer</i>	Specifies the address of the static pointer which was passed to the corresponding call to the <b>mi_open_comm</b> utility to store the address of the module's list of open streams.
<i>Queue</i>	Specifies the <i>Queue</i> parameter which was passed to the corresponding call to the <b>mi_open_comm</b> utility.

## Return Values

The **mi\_close\_comm** utility always returns a value of zero.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts* .

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts* .

The **mi\_open\_comm** utility, **mi\_next\_ptr** utility, **mi\_bufcall** utility.

---

## mi\_next\_ptr Utility

### Purpose

Traverses a STREAMS module's linked list of open streams.

### Syntax

```
#include <pse/mi.h>
#include <sys/stream.h>
```

```
caddr_t mi_next_ptr ( Origin)
caddr_t Origin;
```

### Description

The **mi\_next\_ptr** utility traverses a module's linked list of open streams. The *Origin* argument specifies the address of a per-instance list item, and the return value indicates the address of the next item. The first time the **mi\_next\_ptr** utility is called, the *Origin* parameter should be initialized with the value of the static pointer which was passed to the **mi\_open\_comm** utility. Subsequent calls to the **mi\_next\_ptr** utility should pass the address which was returned by the previous call, until a **NULL** address is returned, indicating that the end of the queue has been reached.

The **mi\_next\_ptr** utility is part of STREAMS kernel extensions.

**Note:** The **stream.h** header file must be the last included header file of each source file using the stream library.

### Parameter

*Origin*        Specifies the address of the current list item being examined.

## Return Values

The **mi\_next\_ptr** utility returns the address of the next list item, or **NULL** if the end of the list has been reached.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts* .

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts* .

The **mi\_close\_comm** utility, **mi\_open\_comm** utility, **mi\_bufcall** utility.

---

## mi\_open\_comm Utility

### Purpose

Performs housekeeping during STREAMS driver or module open operations.

### Syntax

```
#include <pse/mi.h>
#include <sys/stream.h>
```

```
int mi_open_comm ( StaticPointer, Size, Queue, Device, Flag, SFlag, credp)
caddr_t *StaticPointer;
uint Size;
queue_t *Queue;
dev_t *Device;
int Flag;
int SFlag;
cred_t *credp;
```

### Description

The **mi\_open\_comm** subroutine performs housekeeping during STREAMS driver or module open operations. It is intended to be called by the driver or module **open** routine. It assigns a minor device number to the stream (as specified by the *SFlag* parameter), allocates the requested per-stream data, and sets the **q\_ptr** fields of the stream being opened.

The **mi\_open\_comm** subroutine is part of STREAMS kernel extensions.

#### Notes:

1. Each call to the **mi\_open\_comm** subroutine must have a corresponding call to the **mi\_close\_comm** subroutine. Executing one of these utilities without making a corresponding call to the other will lead to unpredictable results.
2. The **stream.h** header file must be the last included header file of each source file using the stream library.

### Parameters

<i>StaticPointer</i>	Specifies the address of a static pointer which will be used internally by the <b>mi_open_comm</b> and related utilities to store the address of the module's list of open streams. This pointer should be initialized to <b>NULL</b> .
<i>Size</i>	Specifies the amount of memory the module needs for its per-stream data. It is usually the size of the local structure which contains the module's instance data.
<i>Queue</i>	Specifies the address of a <b>queue_t</b> structure. The <b>q_ptr</b> field of the of this structure, and of the corresponding read queue structure (if <i>Queue</i> points to a write queue) or write queue structure (if <i>Queue</i> points to a read queue), are filled in with the address of the <b>queue_t</b> structure being initialized.
<i>Device</i>	Specifies the address of a <b>dev_t</b> structure. The use of this parameter depends on the value of the <i>SFlag</i> parameter.
<i>Flag</i>	Unused.

<i>SFlag</i>	Specifies how the <i>Device</i> parameter is to be used. The <i>SFlag</i> parameter may take one of the following values:  <b>DEVOPEN</b> The minor device number specified by the <i>Device</i> argument is used.  <b>MODOPEN</b> The <i>Device</i> parameter is <b>NULL</b> . This value should be used if the <b>mi_open_com</b> subroutine is called from the open routine of a STREAMS module rather than a STREAMS driver.  <b>CLONEOPEN</b> A unique minor device number above 5 is assigned (minor numbers 0-5 are reserved as special access codes).
<i>credp</i>	Unused

## Return Values

On successful completion, the **mi\_open\_comm** subroutine returns a value of zero, otherwise one of the following codes is returned:

Code	Description
<b>ENXIO</b>	Indicates an invalid parameter.
<b>EAGAIN</b>	Indicates that an internal structure could not be allocated, and that the call should be retried.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts* .

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts* .

The **mi\_close\_comm** subroutine, **mi\_next\_ptr** subroutine, **mi\_bufcall** subroutine.

---

## msgdsiz Utility

### Purpose

Gets the number of data bytes in a message.

### Syntax

```
int
msgdsiz(bp)
register mblk_t * bp;
```

### Description

The **msgdsiz** utility returns the number of bytes of data in the message pointed to by the *bp* parameter. Only bytes included in data blocks of type **M\_DATA** are included in the total.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*bp* Specifies the message from which to get the number of bytes.

## Return Values

The **msgdsize** utility returns the number of bytes of data in a message.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## noenable Utility

### Purpose

Prevents a queue from being scheduled.

### Syntax

```
void noenable(q)  
queue_t * q;
```

### Description

The **noenable** utility prevents the queue specified by the *q* parameter from being scheduled for service either by the **putq** or **putbq** utility, when these routines queue an ordinary priority message, or by the **insq** utility when it queues any message. The **noenable** utility does not prevent the scheduling of queues when a high-priority message is queued, unless the message is queued by the **insq** utility.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue to disable.

## Related Information

The **enableok** utility, **insq** utility, **putbq** utility, **putq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## OTHERQ Utility

### Purpose

Returns the pointer to the mate queue.

### Syntax

```
#define OTHERQ( q) ((q)->flag&QREADER? (q)+1: (q)-1)
```

### Description

The **OTHERQ** utility returns a pointer to the mate queue of the *q* parameter.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*q* Specifies that queue whose mate is to be returned.

## Return Values

If the *q* parameter specifies the read queue for the module, the **OTHERQ** utility returns a pointer to the module's write queue. If the *q* parameter specifies the write queue for the module, this utility returns a pointer to the read queue.

## Related Information

The **RD** utility, **WR** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## pfmod Packet Filter Module

### Purpose

Selectively removes upstream data messages on a Stream.

### Synopsis

```
#include <stropts.h>
#include <sys/pfmod.h>

ioctl(fd, I_PUSH, "pfmod");
```

### Description

The pfmod module implements a programmable packet filter facility that may be pushed over any stream. Every data message that pfmod receives on its read side is subjected to a filter program. If the filter program accepts a message, it will be passed along upstream, and will otherwise be freed. If no filter program has been set (as is the case when pfmod is first pushed), all messages are accepted. Non-data messages (for example, M\_FLUSH, M\_PCPROTO, M\_IOCACK) are never examined and always accepted. The write side is not filtered.

Data messages are defined as either M\_PROTO or M\_DATA. If an M\_PROTO message is received, pfmod will skip over all the leading blocks until it finds an M\_DATA block. If none is found, the message is accepted. The M\_DATA portion of the message is then made contiguous with pullupmsg(), if necessary, to ensure the data area referenced by the filter program can be accessed in a single mblk\_t.

### IOCTLs

The following ioctls are defined for this module. All other ioctls are passed downstream without examination.

### PFIOCSETF

Install a new filter program, replacing any previous program. It uses the following data structure:

```
typedef struct packetfilt {
    uchar    Pf_Priority;
    uchar    Pf_FilterLen;
    ushort   Pf_Filter[MAXFILTERS];
} pfilter_t;
```

Pf\_Priority is currently ignored, and should be set to zero. Pf\_FilterLen indicates the number of shortwords in the Pf\_Filter array. Pf\_Filter is an array of shortwords that comprise the filter program. See "Filters" for details on how to write filter programs.

This ioctl may be issued either transparently or as an I\_STR. It will return 0 on success, or -1 on failure, and set errno to one of:

Value	Description
<b>ERANGE</b>	The length of the M_IOCTL message data was not exactly size of (pfilter_t). The data structure is not variable length, although the filter program is.
<b>EFAULT</b>	The ioctl argument points out of bounds.

## Filters

A filter program consists of a linear array of shortword instructions. These instructions operate upon a stack of shortwords. Flow of control is strictly linear; there are no branches or loops. When the filter program completes, the top of the stack is examined. If it is non-zero, or if the stack is empty, the packet being examined is passed upstream (accepted), otherwise the packet is freed (rejected).

Instructions are composed of three portions: push command PF\_CMD(), argument PF\_ARG(), and operation PF\_OP(). Each instruction optionally pushes a shortword onto the stack, then optionally performs a binary operation on the top two elements on the stack, leaving its result on the stack. If there are not at least two elements on the stack, the operation will immediately fail and the packet will be rejected. The argument portion is used only by certain push commands, as documented below.

The following push commands are defined:

Command	Description
<b>PF_NOPUSH</b>	Nothing is pushed onto the stack.
<b>PF_PUSHZERO</b>	Pushes 0x0000.
<b>PF_PUSHONE</b>	Pushes 0x0001.
<b>PF_PUSHFFFF</b>	Pushes 0xffff.
<b>PF_PUSHFF00</b>	Pushes 0xff00.
<b>PF_PUSH00FF</b>	Pushes 0x00ff.
<b>PF_PUSHLIT</b>	Pushes the next shortword in the filter program as literal data. Execution resumes with the next shortword after the literal data.
<b>PF_PUSHWORD+N</b>	Pushes shortword N of the message onto the data stack. N must be in the range 0-255, as enforced by the macro PF_ARG().

The following operations are defined. Each operation pops the top two elements from the stack, and pushes the result of the operation onto the stack. The operations below are described in terms of v1 and v2. The top of stack is popped into v2, then the new top of stack is popped into v1. The result of v1 op v2 is then pushed onto the stack.

Operation	Description
<b>PF_NOP</b>	The stack is unchanged; nothing is popped.
<b>PF_EQ</b>	v1 == v2
<b>PF_NEQ</b>	v1 != v2
<b>PF_LT</b>	v1 < v2
<b>PF_LE</b>	v1 <= v2
<b>PF_GT</b>	v1 > v2
<b>PF_GE</b>	v1 >= v2
<b>PF_AND</b>	v1 & v2; bitwise
<b>PF_OR</b>	v1   v2; bitwise
<b>PF_XOR</b>	v1 ^ v2; bitwise

The remaining operations are "short-circuit" operations. If the condition checked for is found, then the filter program terminates immediately, either accepting or rejecting the packet as specified, without examining the top of stack. If the condition is not found, the filter program continues. These operators do not push any result onto the stack.

Operation	Description
<b>PF_COR</b>	If v1 == v2, accept.
<b>PF_CNOR</b>	If v1 == v2, reject.
<b>PF_CAND</b>	If v1 != v2, reject.
<b>PF_CNAND</b>	If v1 != v2, accept.

If an unknown push command or operation is specified, the filter program terminates immediately and the packet is rejected.

## Configuration

Before using pfmod, it must be loaded into the kernel. This may be accomplished with the **strload** command, using the following syntax:

```
strload -m pfmod
```

This command will load the pfmod into the kernel and make it available to I\_PUSH. Note that attempting to I\_PUSH pfmod before loading it will result in an **EINVAL** error code.

## Example

The following program fragment will push pfmod on a stream, then program it to only accept messages with an Ethertype of 0x8137. This example assumes the stream is a promiscuous DLPI ethernet stream (see **dlpi** for details).

```
#include <stddef.h>
#include <sys/types.h>
#include <netinet/if_ether.h>
#define scale(x) ((x)/sizeof(ushort))
setfilter(int fd)
{
    pfilter_t filter;
    ushort *fp, offset;

    if (ioctl(fd, I_PUSH, "pfmod"))
        return -1;
    offset = scale(offsetof(struct ether_header, ether_type));
    fp = filter.Pf_Filter;

    /* the filter program */
    *fp++ = PF_PUSHLIT;
    *fp++ = 0x8137;
    *fp++ = PF_PUSHWORD + offset;
    *fp++ = PF_EQ;

    filter.Pf_FilterLen = fp - filter.Pf_Filter;

    if (ioctl(fd, PFIOCSETF, &filter))
        return -1;
    return 0;
}
```

This program may be shortened by combining the operation with the push command:

```
*fp++ = PF_PUSHLIT;
*fp++ = 0x8137;
*fp++ = (PF_PUSHWORD + offset) | PF_EQ;
```

The following filter will accept 802.3 frames addressed to either the Netware raw sap 0xff or the 802.2 sap 0xe0:

```
offset = scale(offsetof(struct ie3_hdr, llc));
*fp++ = PF_PUSHWORD + offset; /* get ssap, dsap */
*fp++ = PF_PUSH00FF | PF_AND; /* keep only dsap */
*fp++ = PF_PUSH00FF | PF_COR; /* is dsap == 0xff? */
*fp++ = PF_PUSHWORD + offset; /* get ssap, dsap again */
*fp++ = PF_PUSH00FF | PF_AND; /* keep only dsap */
*fp++ = PF_PUSHLIT | PF_CAND; /* is dsap == 0xe0? */
*fp++ = 0x00e0;
```

Note the use of PF\_COR in this example. If the dsap is 0xff, then the frame is accepted immediately, without continuing the filter program.

---

## pullupmsg Utility

### Purpose

Concatenates and aligns bytes in a message.

### Syntax

```
int
pullupmsg(mp, len)
register struct msgb * mp;
register int len;
```

### Description

The **pullupmsg** utility concatenates and aligns the number of data bytes specified by the *len* parameter of the passed message into a single, contiguous message block. Proper alignment is hardware-dependent. The **pullupmsg** utility only concatenates across message blocks of similar type. It fails if the *mp* parameter points to a message of less than *len* bytes of similar type. If the *len* parameter contains a value of -1, the **pullupmsg** utility concatenates all blocks of the same type at the beginning of the message pointed to by the *mp* parameter.

As a result of the concatenation, the contents of the message pointed to by the *mp* parameter may be altered.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*mp* Specifies the message that is to be aligned.  
*len* Specifies the number of bytes to align.

### Return Values

On success, the **pullupmsg** utility returns a value of 1. On failure, it returns a value of 0.

### Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## putbq Utility

### Purpose

Returns a message to the beginning of a queue.

### Syntax

```
int  
putbq(q, bp)  
register queue_t * q;  
register mblk_t * bp;
```

### Description

The **putbq** utility puts the message pointed to by the *bp* parameter at the beginning of the queue pointed to by the *q* parameter, in a position in accordance with the message type. High-priority messages are placed at the head of the queue, followed by priority-band messages and ordinary messages. Ordinary messages are placed after all high-priority and priority-band messages, but before all other ordinary messages already on the queue. The queue is scheduled in accordance with the rules described in the **putq** utility. This utility is typically used to replace a message on the queue from which it was just removed.

This utility is part of STREAMS Kernel Extensions.

**Note:** A service procedure must never put a high-priority message back on its own queue, as this would result in an infinite loop.

### Parameters

*q* Specifies the queue on which to place the message.  
*bp* Specifies the message to place on the queue.

### Return Values

The **putbq** utility returns a value of 1 on success. Otherwise, it returns a value of 0.

### Related Information

The **putq** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## putctl1 Utility

### Purpose

Passes a control message with a one-byte parameter.

### Syntax

```
int  
putctl1(q, type, param)  
queue_t *q;
```

## Description

The **putctl1** utility creates a control message of the type specified by the *type* parameter with a one-byte parameter specified by the *param* parameter, and calls the put procedure of the queue pointed to by the *q* parameter, with a pointer to the created message as an argument.

The **putctl1** utility allocates new blocks by calling the **allocb** utility.

This utility is part of STREAMS Kernel Extensions.

## Parameters

<i>q</i>	Specifies the queue.
<i>type</i>	Specifies the type of control message.
<i>param</i>	Specifies the one-byte parameter.

## Return Values

On successful completion, the **putctl1** utility returns a value of 1. It returns a value of 0 if it cannot allocate a message block, or if the value of the *type* parameter is **M\_DATA**, **M\_PROTO**, or **M\_PCPROTO**. The **M\_DELAY** type is allowed.

## Related Information

The **allocb** utility, **putctl** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## putctl Utility

### Purpose

Passes a control message.

### Syntax

```
int  
putctl( q,  type)  
queue_t *q;
```

### Description

The **putctl** utility creates a control message of the type specified by the *type* parameter, and calls the put procedure of the queue pointed to by the *q* parameter. The argument of the put procedure is a pointer to the created message. The **putctl** utility allocates new blocks by calling the **allocb** utility.

This utility is part of STREAMS Kernel Extensions.

### Parameters

<i>q</i>	Specifies the queue that contains the desired put procedure.
<i>type</i>	Specifies the type of control message to create.

## Return Values

On successful completion, the **putctl** utility returns a value of 1. It returns a value of 0 if it cannot allocate a message block, or if the value of the *type* parameter is **M\_DATA**, **M\_PROTO**, **M\_PCPROTO**, or **M\_DELAY**.

## Related Information

The **allocb** utility, **putctl1** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## putmsg System Call

### Purpose

Sends a message on a stream.

### Syntax

```
#include <stropts.h>
```

```
int putmsg (fd, ctlptr,
dataptr, flags)
int fd;
struct strbuf * ctlptr;
struct strbuf * dataptr;
int flags;
```

### Description

The **putmsg** system call creates a message from user-specified buffers and sends the message to a STREAMS file. The message may contain either a data part, a control part or both. The data and control parts to be sent are distinguished by placement in separate buffers. The semantics of each part is defined by the STREAMS module that receives the message.

This system call is part of STREAMS Kernel Extensions.

### Parameters

<i>fd</i>	Specifies a file descriptor referencing an open stream.
<i>ctlptr</i>	Holds the control part of the message.
<i>dataptr</i>	Holds the data part of the message.
<i>flags</i>	Indicates the type of message to be sent. Acceptable values are:
<b>0</b>	Sends a nonpriority message.
<b>RS_HIPRI</b>	Sends a priority message.

The *ctlptr* and *dataptr* parameters each point to a **strbuf** structure that contains the following members:

```
int maxlen;    /* not used */
int len;       /* length of data */
char *buf;     /* ptr to buffer */
```

The `len` field in the **strbuf** structure indicates the number of bytes to be sent, and the `buf` field points to the buffer where the control information or data resides. The `maxlen` field is not used in the **putmsg** system call.

To send the data part of a message, the *dataptr* parameter must be nonnull and the `len` field of the *dataptr* parameter must have a value of 0 or greater. To send the control part of a message, the corresponding values must be set for the *ctlptr* parameter. No data (control) part will be sent if either the *dataptr* (*ctlptr*) parameter is null or the `len` field of the *dataptr* (*ctlptr*) parameter is set to -1.

If a control part is specified, and the *flags* parameter is set to **RS\_HIPRI**, a priority message is sent. If the *flags* parameter is set to 0, a nonpriority message is sent. If no control part is specified and the *flags* parameter is set to **RS\_HIPRI**, the **putmsg** system call fails and sets the **errno** global variable to **EINVAL**. If neither a control part nor a data part is specified and the *flags* parameter is set to 0, no message is sent and 0 is returned.

For nonpriority messages, the **putmsg** system call blocks if the stream write queue is full due to internal flow-control conditions. For priority messages, the **putmsg** system call does not block on this condition. For nonpriority messages, the **putmsg** system call does not block when the write queue is full and the **O\_NDELAY** or **O\_NONBLOCK** flag is set. Instead, the system call fails and sets the **errno** global variable to **EAGAIN**.

The **putmsg** system call also blocks, unless prevented by lack of internal resources, while waiting for the availability of message blocks in the stream, regardless of priority or whether the **O\_NDELAY** or **O\_NONBLOCK** flag has been specified. No partial message is sent.

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and the **errno** global variable is set to indicate the error.

## Error Codes

The **putmsg** system call fails if one of the following is true:

Value	Description
<b>EAGAIN</b>	A nonpriority message was specified, the <b>O_NONBLOCK</b> flag is set, and the stream write queue is full due to internal flow-control conditions.
<b>EAGAIN</b>	Buffers could not be allocated for the message that was to be created.
<b>EBADF</b>	The value of the <i>fd</i> parameter is not a valid file descriptor open for writing.
<b>EFAULT</b>	The <i>ctlptr</i> or <i>dataptr</i> parameter points outside the allocated address space.
<b>EINTR</b>	A signal was caught during the <b>putmsg</b> system call.
<b>EINVAL</b>	An undefined value was specified in the <i>flags</i> parameter, or the <i>flags</i> parameter is set to <b>RS_HIPRI</b> and no control part was supplied.
<b>EINVAL</b>	The stream referenced by the <i>fd</i> parameter is linked below a multiplexer.
<b>ENOSR</b>	Buffers could not be allocated for the message that was to be created due to insufficient STREAMS memory resources.
<b>ENOSTR</b>	A stream is not associated with the <i>fd</i> parameter.
<b>ENXIO</b>	A hangup condition was generated downstream for the specified stream.
<b>EPIPE</b> or <b>EIO</b>	The <i>fd</i> parameter refers to a STREAM-based pipe and the other end of the pipe is closed. A SIGPIPE signal is generated for the calling thread.
<b>ERANGE</b>	The size of the data part of the message does not fall within the range specified by the maximum and minimum packet sizes of the topmost STREAMS module. OR The control part of the message is larger than the maximum configured size of the control part of a message. OR

Value	Description
	The data part of a message is larger than the maximum configured size of the data part of a message.

The **putmsg** system call also fails if a STREAMS error message was processed by the stream head before the call. The error returned is the value contained in the STREAMS error message.

## Files

**/lib/pse.exp** Contains the STREAMS export symbols.

## Related Information

The **getmsg** system call, **getpmsg** system call, **putpmsg** system call.

The **read** subroutine, **poll** subroutine, **write** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## putnext Utility

### Purpose

Passes a message to the next queue.

### Syntax

```
#define putnext( q, mp) ((*q)->q_next->q_qinfo->q_i_putp)((q)-q_next, (mp))
```

### Description

The **putnext** utility calls the put procedure of the next queue in a stream and passes to the procedure a message pointer as an argument. The **putnext** utility is the typical means of passing messages to the next queue in a stream.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the calling queue.  
*mp* Specifies the message that is to be passed.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## putpmsg System Call

### Purpose

Sends a priority message on a stream.

## Syntax

```
#include <stropts.h>
```

```
int putpmsg (fd, ctlptr,  
dataptr, band, flags)  
int fd;  
struct strbuf * ctlptr;  
struct strbuf * dataptr;  
int band;  
int flags;
```

## Description

The **putpmsg** system call is identical to the **putmsg** system call except that it sends a priority message. All information except for flag settings are found in the description for the **putmsg** system call. The differences in the flag settings are noted in the error codes section.

This system call is part of STREAMS Kernel Extensions.

## Parameters

<i>fd</i>	Specifies a file descriptor referencing an open stream.
<i>ctlptr</i>	Holds the control part of the message.
<i>dataptr</i>	Holds the data part of the message.
<i>band</i>	Indicates the priority band.
<i>flags</i>	Indicates the priority type of message to be sent. Acceptable values are:

### **MSG\_BAND**

Sends a non-priority message.

### **MSG\_HIPRI**

Sends a priority message.

## Error Codes

The **putpmsg** system call is unsuccessful under the following conditions:

- The *flags* parameter is set to a value of 0.
- The *flags* parameter is set to **MSG\_HIPRI** and the *band* parameter is set to a nonzero value.
- The *flags* parameter is set to **MSG\_HIPRI** and no control part is specified.

## Related Information

The **poll** subroutine, **read** subroutine, **write** subroutine.

The **getmsg** system call, **getpmsg** system call, **putmsg** system call.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## putq Utility

### Purpose

Puts a message on a queue.

## Syntax

```
int  
putq(q, bp)  
register queue_t * q;  
register mblk_t * bp;
```

## Description

The **putq** utility puts the message pointed to by the *bp* parameter on the message queue pointed to by the *q* parameter, and then enables that queue. The **putq** utility queues messages based on message-queuing priority.

The priority classes are:

Class	Description
type >= QPCTL	High-priority
type < QPCTL && band > 0	Priority band
type < QPCTL && band == 0	Normal

When a high-priority message is queued, the **putq** utility always enables the queue. For a priority-band message, the **putq** utility is allowed to enable the queue (the **QNOENAB** flag is not set). Otherwise, the **QWANTR** flag is set, indicating that the service procedure is ready to read the queue. When an ordinary message is queued, the **putq** utility enables the queue if the following condition is set and if enabling is not inhibited by the **noenable** utility: the module has just been pushed, or else no message was queued on the last **getq** call and no message has been queued since.

The **putq** utility looks only at the priority band in the first message block of a message. If a high-priority message is passed to the **putq** utility with a nonzero *b\_band* field value, the *b\_band* field is reset to 0 before the message is placed on the queue. If the message passed to the **putq** utility has a *b\_band* field value greater than the number of **qband** structures associated with the queue, the **putq** utility tries to allocate a new **qband** structure for each band up to and including the band of the message.

The **putq** utility should be used in the put procedure for the same queue in which the message is queued. A module should not call the **putq** utility directly in order to pass messages to a neighboring module. Instead, the **putq** utility itself can be used as the value of the *qi\_putp* field in the put procedure for either or both of the module **qinit** structures. Doing so effectively bypasses any put-procedure processing and uses only the module service procedures.

This utility is part of STREAMS Kernel Extensions.

**Note:** The service procedure must never put a priority message back on its own queue, as this would result in an infinite loop.

## Parameters

*q* Specifies the queue on which to place the message.  
*bp* Specifies the message to put on the queue.

## Return Values

On successful completion, the **putq** utility returns a value of 1. Otherwise, it returns a value of 0.

## Related Information

The **getq** utility.

## qenable Utility

### Purpose

Enables a queue.

### Syntax

```
void qenable (q)
register queue_t * q;
```

### Description

The **qenable** utility places the queue pointed to by the *q* parameter on the linked list of queues ready to be called by the STREAMS scheduler.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the queue to be enabled.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## qreply Utility

### Purpose

Sends a message on a stream in the reverse direction.

### Syntax

```
void qreply (q, bp)
register queue_t * q;
register mblk_t * bp;
```

### Description

The **qreply** utility sends the message pointed to by the *bp* parameter either up or down the stream-in the reverse direction from the queue pointed to by the *q* parameter. The utility does this by locating the partner of the queue specified by the *q* parameter (see the **OTHERQ** utility), and then calling the put procedure of that queue's neighbor (as in the **putnext** utility). The **qreply** utility is typically used to send back a response (**M\_IOCACK** or **M\_IOCNAK** message) to an **M\_IOCTL** message.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies which queue to send the message up or down.  
*bp* Specifies the message to send.

## Related Information

The **OTHERQ** utility, **putnext** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## qsize Utility

### Purpose

Finds the number of messages on a queue.

### Syntax

```
int
qsize(qp)
register queue_t * qp;
```

### Description

The **qsize** utility returns the number of messages present in the queue specified by the *qp* parameter. If there are no messages on the queue, the **qsize** parameter returns a value of 0.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*qp* Specifies the queue on which to count the messages.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## RD Utility

### Purpose

Gets the pointer to the read queue.

### Syntax

```
#define RD( q) ((q)-1)
```

### Description

The **RD** utility accepts a write-queue pointer, specified by the *q* parameter, as an argument and returns a pointer to the read queue for the same module.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the write queue.

## Related Information

The **OTHERQ** utility, **WR** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## rmvb Utility

### Purpose

Removes a message block from a message.

### Syntax

```
mb1k_t *  
rmvb(mp, bp)  
register mb1k_t * mp;  
register mb1k_t * bp;
```

### Description

The **rmvb** utility removes the message block pointed to by the *bp* parameter from the message pointed to by the *mp* parameter, and then restores the linkage of the message blocks remaining in the message. The **rmvb** utility does not free the removed message block, but returns a pointer to the head of the resulting message. If the message block specified by the *bp* parameter is not contained in the message specified by the *mp* parameter, the **rmvb** utility returns a -1. If there are no message blocks in the resulting message, the **rmvb** utility returns a null pointer.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*bp* Specifies the message block to be removed.  
*mp* Specifies the message from which to remove the message block.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## rmvq Utility

### Purpose

Removes a message from a queue.

### Syntax

```
void rmvq(q, mp)  
register queue_t * q;  
register mb1k_t * mp;
```

### Description

**Attention:** If the *mp* parameter does not point to a message that is present on the specified queue, a system panic could result.

The **rmvq** utility removes the message pointed to by the *mp* parameter from the message queue pointed to by the *q* parameter, and then restores the linkage of the messages remaining on the queue.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*q* Specifies the queue from which to remove the message.  
*mp* Specifies the message to be removed.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## sad Device Driver

### Purpose

Provides an interface for administrative operations.

### Syntax

```
#include <sys/types.h>
#include <sys/conf.h>
#include <sys/sad.h>
#include <sys/stropts.h>
```

```
int ioctl (fildes, command, arg)
int fildes, command;
int arg;
```

### Description

The STREAMS Administrative Driver (**sad**) provides an interface for applications to perform administrative operations on STREAMS modules and drivers. The interface is provided through **ioctl** operations. Privileged operation can access the **sad** device driver in the **/dev/sad/user** directory.

### Parameters

*fildes* Specifies an open file descriptor that refers to the **sad** device driver.  
*command* Determines the control function to be performed.  
*arg* Supplies additional information for the given control function.

### Values for the command Parameter

The **autopush** command allows a user to configure a list of modules to be automatically pushed on a stream when a driver is first opened. The **autopush** command is controlled by the following commands.

Command	Description
<b>SAD_SAP</b>	<p>Allows the person performing administrative duties to configure the information for the given device, which is used by the <b>autopush</b> command. The <i>arg</i> parameter points to a <b>strapush</b> structure containing the following elements:</p> <pre>uint sap_cmd; long sap_major; long sap_minor; long sap_lastminor; long sap_npush; uint sap_list[MAXAPUSH] [FMNAMESZ + 1];</pre> <p>The elements are described as follows:</p> <p><b>sap_cmd</b> Indicates the type of configuration being done. Acceptable values are:</p> <p><b>SAP_ONE</b> Configures one minor device of a driver.</p> <p><b>SAP_RANGE</b> Configures a range of minor devices of a driver.</p> <p><b>SAP_ALL</b> Configures all minor devices of a driver.</p> <p><b>SAP_CLEAR</b> Undoes configuration information for a driver.</p> <p><b>sap_major</b> Specifies the major device number of the device to be configured.</p> <p><b>sap_minor</b> Specifies the minor device number of the device to be configured.</p> <p><b>sap_lastminor</b> Specifies the last minor device number in a range of devices to be configured. This field is used only with the <b>SAP_RANGE</b> value in the <i>sap_cmd</i> field.</p> <p><b>sap_npush</b> Indicates the number of modules to be automatically pushed when the device is opened. The value of this field must be less than or equal to <b>MAXAPUSH</b>, which is defined in the <b>sad.h</b> file. It must also be less than or equal to <b>NSTRPUSH</b>, which is defined in the kernel master file.</p> <p><b>sap_list</b> Specifies an array of module names to be pushed in the order in which they appear in the list.</p> <p>When using the <b>SAP_CLEAR</b> value, the user sets only the <i>sap_major</i> and <i>sap_minor</i> fields. This undoes the configuration information for any of the other values. If a previous entry was configured with the <b>SAP_ALL</b> value, the <i>sap_minor</i> field is set to 0. If a previous entry was configured with the <b>SAP_RANGE</b> value, the <i>sap_minor</i> field is set to the lowest minor device number in the range configured.</p> <p>On successful completion, the return value from the <b>ioctl</b> operation is 0. Otherwise, the return value is -1.</p>
<b>SAD_GAP</b>	<p>Allows any user to query the <b>sad</b> device driver to get the <b>autopush</b> configuration information for a given device. The <i>arg</i> parameter points to a <b>strapush</b> structure as described under the <b>SAD_SAP</b> value.</p> <p>The user sets the <i>sap_major</i> and <i>sap_minor</i> fields to the major and minor device numbers, respectively, of the device in question. On return, the <b>strapush</b> structure is filled with the entire information used to configure the device. Unused entries are filled with zeros.</p> <p>On successful completion, the return value from the <b>ioctl</b> operation is 0. Otherwise, the return value is -1.</p>

Command	Description
<b>SAD_VML</b>	<p>Allows any user to validate a list of modules; that is, to see if they are installed on the system. The <i>arg</i> parameter is a pointer to a <b>str_list</b> structure containing the following elements:</p> <pre>int sl_nmods; struct str_mlist *sl_modlist;</pre> <p>The <b>str_mlist</b> structure contains the following element:</p> <pre>char l_name[FMNAMESZ+1];</pre> <p>The fields are defined as follows:</p> <p><b>sl_nmods</b> Indicates the number of entries the user has allocated in the array.</p> <p><b>sl_modlist</b> Points to the array of module names.</p>

## Return Values

On successful completion, the return value from the **ioctl** operation is 0 if the list is valid or 1 if the list contains an invalid module name. Otherwise the return value is -1.

## Error Codes

On failure, the **errno** global variable is set to one of the following values:

Value	Description
<b>EFAULT</b>	The <i>arg</i> parameter points outside the allocated address space.
<b>EINVAL</b>	The major device number is not valid, the number of modules is not valid.
	OR
	The list of module names is not valid.
<b>ENOSTR</b>	The major device number does not represent a STREAMS driver.
<b>EEXIST</b>	The major-minor device pair is already configured.
<b>ERANGE</b>	The value of the <i>command</i> parameter is <b>SAP_RANGE</b> and the value in the <i>sap_lastminor</i> field is not greater than the value in the <i>sap_minor</i> field.
	OR
	The value of the <i>command</i> parameter is <b>SAP_CLEAR</b> and the value in the <i>sap_minor</i> field is not equal to the first minor in the range.
<b>ENODEV</b>	The value in the <i>command</i> parameter is <b>SAP_CLEAR</b> and the device is not configured for the <b>autopush</b> command.
<b>ENOSR</b>	An internal <b>autopush</b> data structure cannot be allocated.

## Related Information

The **autopush** command.

The **close** subroutine, **fstat** subroutine, **open** subroutine, **stat** subroutine.

Understanding streamio (STREAMS ioctl) Operations, Understanding STREAMS Drivers and Modules, Understanding the log Device Driver in *AIX Version 6.1 Communications Programming Concepts*.

---

## splstr Utility

### Purpose

Sets the processor level.

### Syntax

```
int splstr()
```

### Description

The **splstr** utility increases the system processor level in order to block interrupts at a level appropriate for STREAMS modules and drivers when they are executing critical portions of their code. The **splstr** utility returns the processor level at the time of its invocation. Module developers are expected to use the standard **splx(s)** utility, where **s** is the integer value returned by the **splstr** operation, to restore the processor level to its previous value after the critical portions of code are passed.

This utility is part of STREAMS Kernel Extensions.

### Related Information

The **splx** utility.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## splx Utility

### Purpose

Terminates a section of code.

### Syntax

```
int splx(x)
int x;
```

### Description

The **splx** utility terminates a section of protected critical code. This utility restores the interrupt level to the previous level specified by the *x* parameter.

This utility is part of STREAMS Kernel Extensions.

### Related Information

The **splstr** utility.

List of Streams Programming References and Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

---

## srv Utility

### Purpose

Services queued messages for STREAMS modules or drivers.

## Syntax

```
#include <sys/types.h>
#include <sys/stream.h>
#include <sys/stropts.h>

int <prefix>rsrv(queue_t *q);           /* read side */
int <prefix>wsrv(queue_t *q);           /* write side */
```

## Parameters

*q*      Pointer to the queue structure.

## Description

The optional service (<prefix>**srv**) routine can be included in a STREAMS module or driver for one or more of the following reasons:

- To provide greater control over the flow of messages in a stream
- To make it possible to defer the processing of some messages to avoid depleting system resources
- To combine small messages into larger ones, or break large messages into smaller ones
- To recover from resource allocation failure. A module's or driver's **put** routine can test for the availability of a resource, and if it is not available, enqueue the message for later processing by the **srv** routine.

A message is first passed to a module's or driver's **put** routine, which may or may not do some processing. It must then either:

- Pass the message to the next stream component with **putnext**
- If a **srv** routine has been included, it may call **putq** to place the message on the queue

Once a message has been enqueued, the STREAMS scheduler controls the invocation of the service routine. Service routines are called in FIFO order by the scheduler. No guarantees can be made about how long it will take for a **srv** routine to be called except that it will happen before any user level process is run.

Every stream component (stream head, module or driver) has limit values it uses to implement flow control. Tunable high and low water marks should be checked to stop and restart the flow of message processing. Flow control limits apply only between two adjacent components with **srv** routines.

STREAMS messages can be defined to have up to 256 different priorities to support requirements for multiple bands of data flow. At a minimum, a stream must distinguish between normal (priority zero) messages and high priority messages (such as M\_IOCACK). High priority messages are always placed at the head of the **srv** routine's queue, after any other enqueued high priority messages. Next are messages from all included priority bands, which are enqueued in decreasing order of priority. Each priority band has its own flow control limits. If a flow controlled band is stopped, all lower priority bands are also stopped.

Once the STREAMS scheduler calls a **srv** routine, it must process all messages on its queue. The following steps are general guidelines for processing messages. Keep in mind that many of the details of how a **srv** routine should be written depend on the implementation, the direction of flow (upstream or downstream), and whether it is for a module or a driver.

1. Use **getq** to get the next enqueued message.
2. If the message is high priority, process it (if appropriate) and pass it to the next stream component with **putnext**.
3. If it is not a high priority message (and therefore subject to flow control), attempt to send it to the next stream component with a **srv** routine. Use **canput** or **bcanput** to determine if this can be done.

4. If the message cannot be passed, put it back on the queue with **putbq**. If it can be passed, process it (if appropriate) and pass it with **putnext**.

Rules for service routines:

1. Service routines must not call any kernel services that sleep or are not interrupt safe.
2. Service routines are called by the STREAMS scheduler with most interrupts enabled.

**Note:** Each stream module must specify a read and a write service (**srv**) routine. If a service routine is not needed (because the **put** routine processes all messages), a NULL pointer should be placed in module's qinit structure. Do not use **nulldev** instead of the NULL pointer. Use of **nulldev** for a **srv** routine may result in flow control errors.

Prior to AIX 4.1, STREAMS service routines were permitted which were not coded to specification (that is, the service routine called sleep or called kernel services that slept, other possibilities). In AIX 4.1, this behavior will cause a system failure because the STREAMS scheduler is executed with some interrupts disabled. Modules or drivers can force the old style scheduling by setting the `sc_flags` field of the `kstrconf_t` structure to `STR_Q_NOTTOSPEC`. This structure is passed to the system when the module or driver calls the `str_install` STREAMS service. This flag will cause STREAMS to schedule the module's or driver's service routines with all interrupts enabled. There is a severe performance penalty for this type of STREAMS scheduling and future releases may not support `STR_Q_NOTTOSPEC`.

## Return Values

Ignored.

## Related Information

**put**, **bcanput**, **canput**, **getq**, **putbq**, **putnext**, **putq** utilities.

The **queue** structure in `/usr/include/sys/stream.h`.

---

## str\_install Utility

### Purpose

Installs streams modules and drivers.

### Syntax

```
#include <sys/strconf.h>
```

```
int
str_install(cmd, conf)
int cmd;
strconf_t * conf;
```

### Description

The **str\_install** utility adds or removes Portable Streams Environment (PSE) drivers and modules from the internal tables of PSE. The extension is pinned when added and unpinned when removed (see the **pincode** kernel service). It uses a configuration structure to provide sufficient information to perform the specified command.

This utility is part of STREAMS Kernel Extensions.

The configuration structure, **strconf\_t**, is defined as follows:

```
typedef struct {
    char *sc_name;
    struct streamtab *sc_str;
    int sc_open_style;
    int sc_flags;
    int sc_major;
    int sc_sqllevel;
    caddr_t sc_sqinfo;
} strconf_t;
```

The elements of the **strconf\_t** structure are defined as follows:

Element	Description
sc_name	Specifies the name of the extension in the internal tables of PSE. For modules, this name is installed in the <b>fmodsw</b> table and is used for <b>I_PUSH</b> operations. For drivers, this name is used only for reporting with the <b>scls</b> and <b>strinfo</b> commands.
sc_str	Points to a <b>streamtab</b> structure.

**Element**  
sc\_open\_style sc\_flags

**Description**

Specifies the style of the driver or module open routine. The acceptable values are:

**STR\_NEW\_OPEN**

Specifies the open syntax and semantics used in System V Release 4.

**STR\_OLD\_OPEN**

Specifies the open syntax and semantics used in System V Release 3.

If the module is multiprocessor-safe, the following flag should be added by using the bitwise OR operator:

**STR\_MPSAFE**

Specifies that the extension was designed to run on a multiprocessor system.

If the module uses callback functions that need to be protected against interrupts (non-interrupt-safe callback functions) for the **timeout** or **bufcall** utilities, the following flag should be added by using the bitwise OR operator:

**STR\_QSAFETY**

Specifies that the extension uses non-interrupt-safe callback functions for the **timeout** or **bufcall** utilities.

This flag is automatically set by STREAMS if the module is not multiprocessor-safe.

**STR\_PERSTREAM**

Specifies that the module accepts to run at perstream synchronization level.

**STR\_Q\_NOTTOSPEC**

Specifies that the extension is designed to run its service routine under process context.

By default STREAMS service routine runs under interrupt context (INTOFFL3). If Streams drivers or modules want to execute their service routine under process context (INTBASE), they need to set this flag.

**STR\_64BIT**

Specifies that the extension is capable to support 64-bit data types.

**STR\_NEWCLONING**

Specifies the driver open uses new-style cloning. Under this style, the driver open() is not checking for CLONEOPEN flag and returns new device number.

sc\_major

Specifies the major number of the device.

Element	Description
sc_sqllevel	<p>Reserved for future use. Specifies the synchronization level to be used by PSE. There are seven levels of synchronization:</p> <p><b>SQLVL_NOP No synchronization</b> Specifies that each queue can be accessed by more than one thread at the same time. The protection of internal data and of <b>put</b> and <b>service</b> routines against the <b>timeout</b> or <b>bufcall</b> utilities is done by the module or driver itself. This synchronization level should be used essentially for multiprocessor-efficient modules.</p> <p><b>SQLVL_QUEUE Queue Level</b> Specifies that each queue can be accessed by only one thread at the same time. This is the finest synchronization level, and should only be used when the two sides of a queue pair do not share common data.</p> <p><b>SQLVL_QUEUEPAIR Queue Pair Level</b> Specifies that each queue pair can be accessed by only one thread at the same time.</p> <p><b>SQLVL_MODULE Module Level</b> Specifies that all instances of a module can be accessed by only one thread at the same time. This is the default value.</p> <p><b>SQLVL_ELSEWHERE Arbitrary Level</b> Specifies that a group of modules can be accessed by only one thread at the same time. Usually, the group of modules is a set of cooperating modules, such as a protocol family. The group is defined by using the same name in the <code>sc_sqinfo</code> field for each module in the group.</p> <p><b>SQLVL_GLOBAL Global Level</b> Specifies that all of PSE can be accessed by only one thread at the same time. This option should normally be used only for debugging.</p> <p><b>SQLVL_DEFAULT Default Level</b> Specifies the default level, set to <b>SQLVL_MODULE</b>.</p>
sc_sqinfo	<p>Specifies an optional group name. This field is only used when the <b>SQLVL_ELSEWHERE</b> arbitrary synchronization level is set; all modules having the same name belong to one group. The name size is limited to eight characters.</p>

## Parameters

<i>cmd</i>	<p>Specifies which operation to perform. Acceptable values are:</p> <p><b>STR_LOAD_DEV</b> Adds a device into PSE internal tables.</p> <p><b>STR_UNLOAD_DEV</b> Removes a device from PSE internal tables.</p> <p><b>STR_LOAD_MOD</b> Adds a module into PSE internal tables.</p> <p><b>STR_UNLOAD_MOD</b> Removes a module from PSE internal tables.</p>
<i>conf</i>	<p>Points to a <b>strconf_t</b> structure, which contains all the necessary information to successfully load and unload a PSE kernel extension.</p>

## Return Values

On successful completion, the **str\_install** utility returns a value of 0. Otherwise, it returns an error code.

## Error Codes

On failure, the **str\_install** utility returns one of the following error codes:

Code	Description
<b>EBUSY</b>	The PSE kernel extension is already in use and cannot be unloaded.
<b>EEXIST</b>	The PSE kernel extension already exists in the system.
<b>EINVAL</b>	A parameter contains an unacceptable value.
<b>ENODEV</b>	The PSE kernel extension could not be loaded.
<b>ENOENT</b>	The PSE kernel is not present and could not be unloaded.
<b>ENOMEM</b>	Not enough memory for the extension could be allocated and pinned.
<b>ENXIO</b>	PSE is currently locked for use.

## Related Information

The **pincode** kernel service, **unpincode** kernel service.

The **streamio** operations.

Configuring Drivers and Modules in the Portable Streams Environment (PSE) and List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts*.

---

## streamio Operations

### Purpose

Perform a variety of control functions on streams.

### Syntax

```
#include <stropts.h>
int ioctl (fildes, command, arg)
int fildes, command;
```

### Description

See individual **streamio** operations for a description of each one.

This operation is part of STREAMS Kernel Extensions.

### Parameters

<i>fildes</i>	Specifies an open file descriptor that refers to a stream.
<i>command</i>	Determines the control function to be performed.
<i>arg</i>	Represents additional information that is needed by this operation.

The type of the *arg* parameter depends upon the operation, but it is generally an integer or a pointer to a *command*-specific data structure.

The *command* and *arg* parameters are passed to the file designated by the *fildes* parameter and are interpreted by the stream head. Certain combinations of these arguments can be passed to a module or driver in the stream.

## Values of the command Parameter

The following ioctl operations are applicable to all STREAMS files:

Operation	Description
<b>I_ATMARK</b>	Checks if the current message on the stream-head read queue is marked.
<b>I_CANPUT</b>	Checks if a given band is writable.
<b>I_CKBAND</b>	Checks if a message of a particular band is on the stream-head queue.
<b>I_FDINSERT</b>	Creates a message from user specified buffers, adds information about another stream and sends the message downstream.
<b>I_FIND</b>	Compares the names of all modules currently present in the stream to a specified name.
<b>I_FLUSH</b>	Flushes all input or output queues.
<b>I_FLUSHBAND</b>	Flushes all message of a particular band.
<b>I_GETBAND</b>	Gets the band of the first message on the stream-head read queue.
<b>I_GETCLTIME</b>	Returns the delay time.
<b>I_GETSIG</b>	Returns the events for which the calling process is currently registered to be sent a <b>SIGPOLL</b> signal.
<b>I_GRDOPT</b>	Returns the current read mode setting.
<b>I_LINK</b>	Connects two specified streams.
<b>I_LIST</b>	Lists all the module names on the stream.
<b>I_LOOK</b>	Retrieves the name of the module just below the stream head.
<b>I_NREAD</b>	Counts the number of data bytes in data blocks in the first message on the stream-head read queue, and places this value in a specified location.
<b>I_PEEK</b>	Allows a user to retrieve the information in the first message on the stream-head read queue without taking the message off the queue.
<b>I_PLINK</b>	Connects two specified streams.
<b>I_POP</b>	Removes the module just below the stream head.
<b>I_PUNLINK</b>	Disconnects the two specified streams.
<b>I_PUSH</b>	Pushes a module onto the top of the current stream.
<b>I_RECVFD</b>	Retrieves the file descriptor associated with the message sent by an <b>I_SENDFD</b> operation over a stream pipe.
<b>I_SENDFD</b>	Requests a stream to send a message to the stream head at the other end of a stream pipe.
<b>I_SETCLTIME</b>	Sets the time that the stream head delays when a stream is closing.
<b>I_SETSIG</b>	Informs the stream head that the user wishes the kernel to issue the <b>SIGPOLL</b> signal when a particular event occurs on the stream.
<b>I_SRDOPT</b>	Sets the read mode.
<b>I_STR</b>	Constructs an internal STREAMS ioctl message.

Operation	Description
<b>I_UNLINK</b>	Disconnects the two specified streams.

## Return Values

Unless specified otherwise, the return value from the **ioctl** subroutine is 0 upon success and -1 if unsuccessful with the **errno** global variable set as indicated.

## Related Information

List of Streams Programming References, Understanding streamio (STREAMS ioctl) Operations in *AIX Version 6.1 Communications Programming Concepts*.

---

## strlog Utility

### Purpose

Generates STREAMS error-logging and event-tracing messages.

### Syntax

```
int
strlog(mid, sid, level, flags, fmt, arg1, . . . )
```

```
short mid, sid;
char level;
ushort flags;
char * fmt;
unsigned arg1;
```

### Description

The **strlog** utility generates log messages within the kernel. Required definitions are contained in the **sys/strlog.h** file.

This utility is part of STREAMS Kernel Extensions.

### Parameters

<i>mid</i>	Specifies the STREAMS module ID number for the module or driver submitting the <b>log</b> message.
<i>sid</i>	Specifies an internal sub-ID number usually used to identify a particular minor device of a driver.
<i>level</i>	Specifies a tracing level that allows for selective screening of low-priority messages from the tracer.

<i>flags</i>	Specifies the destination of the message. This can be any combination of:  <b>SL_ERROR</b> The message is for the error logger.  <b>SL_TRACE</b> The message is for the tracer.  <b>SL_CONSOLE</b> Log the message to the console.  <b>SL_FATAL</b> Advisory notification of a fatal error.  <b>SL_WARN</b> Advisory notification of a nonfatal error.  <b>SL_NOTE</b> Advisory message.  <b>SL_NOTIFY</b> Request that a copy of the message be mailed to the system administrator.
<i>fmt</i>	Specifies a print style-format string, except that %f, %e, %E, %g, and %G conversion specifications are not handled.
<i>arg1</i>	Specifies numeric or character arguments. Up to <b>NLOGARGS</b> (currently 4) numeric or character arguments can be provided. (The <b>NLOGARGS</b> variable specifies the maximum number of arguments allowed. It is defined in the <b>sys/strlog.h</b> file.)

## Related Information

The **streamio** operations.

**clone** Device Driver in *AIX Version 6.1 Communications Programming Concepts*.

List of Streams Programming References, Understanding the log Device Driver, Understanding STREAMS Error and Trace Logging in *AIX Version 6.1 Communications Programming Concepts*.

---

## strqget Utility

### Purpose

Obtains information about a queue or band of the queue.

### Syntax

```
int
strqget(q, what, pri, valp)
register queue_t * q;
qfields_t what;
register unsigned char pri;
long * valp;
```

### Description

The **strqget** utility allows modules and drivers to get information about a queue or particular band of the queue. The information is returned in the *valp* parameter. The fields that can be obtained are defined as follows:

This utility is part of STREAMS Kernel Extensions.

```
typedef enum qfields {
    QHIWAT    = 0,
    QLOWAT    = 1,
    QMAXPSZ   = 2,
    QMINPSZ   = 3,
    QCOUNT   = 4,
    QFIRST    = 5,
    QLAST     = 6,
    QFLAG     = 7,
    QBAD      = 8
} qfields_t;
```

## Parameters

*q* Specifies the queue about which to get information.  
*what* Specifies the information to get from the queue.  
*pri* Specifies the priority band about which to get information.  
*valp* Contains the requested information on return.

## Return Values

On success, the **strqget** utility returns a value of 0. Otherwise, it returns an error number.

## Related Information

List of Streams Programming References, Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_accept Subroutine for Transport Layer Interface

### Purpose

Accepts a connect request.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_accept (fd, resfd, call)
int fd;
int resfd;
struct t_call * call;
```

### Description

The **t\_accept** subroutine is issued by a transport user to accept a connect request. A transport user can accept a connection on either the same local transport end point or on an end point different from the one on which the connect indication arrived.

### Parameters

*fd* Identifies the local transport end point where the connect indication arrived.  
*resfd* Specifies the local transport end point where the connection is to be established.

*call* Contains information required by the transport provider to complete the connection. The *call* parameter points to a **t\_call** structure, which contains the following members:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;  
int sequence;
```

The **netbuf** structure is described in the **tiuser.h** file. In the *call* parameter, the *addr* field is the address of the caller, the *opt* field indicates any protocol-specific parameters associated with the connection, the *udata* field points to any user data to be returned to the caller, and the *sequence* field is the value returned by the **t\_listen** subroutine which uniquely associates the response with a previously received connect indication.

If the same end point is specified (that is, the *resfd* value equals the *fd* value), the connection can be accepted unless the following condition is true: the user has received other indications on that end point, but has not responded to them (with either the **t\_accept** or **t\_snddis** subroutine). For this condition, the **t\_accept** subroutine fails and sets the **t\_errno** variable to **TBADF**.

If a different transport end point is specified (that is, the *resfd* value does not equal the *fd* value), the end point must be bound to a protocol address and must be in the **T\_IDLE** state (see the **t\_getstate** subroutine) before the **t\_accept** subroutine is issued.

For both types of end points, the **t\_accept** subroutine fails and sets the **t\_errno** variable to **TLOOK** if there are indications (for example, a connect or disconnect) waiting to be received on that end point.

The values of parameters specified by the *opt* field and the syntax of those values are protocol-specific. The *udata* field enables the called transport user to send user data to the caller, the amount of user data must not exceed the limits supported by the transport provider as returned by the **t\_open** or **t\_getinfo** subroutine. If the value in the *len* field of the *udata* field is 0, no data will be sent to the caller.

## Return Values

On successful completion, the **t\_connect** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to accept a connection on the responding transport end point or use the specified options.
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADDR</b>	The specified protocol address was in an incorrect format or contained illegal information.
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point; or the user is illegally accepting a connection on the same transport end point on which the connect indication arrived.
<b>TBADOPT</b>	The specified options were in an incorrect format or contained illegal information.
<b>TBADSEQ</b>	An incorrect sequence number was specified.
<b>TLOOK</b>	An asynchronous event has occurred on the transport end point referenced by the <i>fd</i> parameter and requires immediate attention.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The function was issued in the wrong sequence on the transport end point referenced by the <i>fd</i> parameter, or the transport end point referred to by the <i>resfd</i> parameter is not in the <b>T_IDLE</b> state.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_alloc** subroutine, **t\_connect** subroutine, **t\_getinfo** subroutine, **t\_getstate** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_rcvconnect** subroutine and **t\_snddis** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_alloc Subroutine for Transport Layer Interface

### Purpose

Allocates a library structure.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
char *t_alloc (fd, struct_type, fields)
int fd;
int struct_type;
int fields;
```

### Description

The **t\_alloc** subroutine dynamically assigns memory for the various transport-function argument structures. This subroutine allocates memory for the specified structure, and also allocates memory for buffers referenced by the structure.

Use of the **t\_alloc** subroutine to allocate structures will help ensure the compatibility of user programs with future releases of the transport interface.

### Parameters

*fd* Specifies the transport end point through which the newly allocated structure will be passed.

*struct\_type* Specifies the structure to be allocated. The structure to allocate is specified by the *struct\_type* parameter, and can be one of the following:

**T\_BIND**  
     **struct t\_bind**

**T\_CALL**  
     **struct t\_call**

**T\_OPTMGMT**  
     **struct t\_optmgmt**

**T\_DIS**   **struct t\_discon**

**T\_UNITDATA**  
     **struct t\_unitdata**

**T\_UDERROR**  
     **struct t\_uderr**

**T\_INFO**  
     **struct t\_info**

Each of these structures may subsequently be used as a parameter to one or more transport functions.

Each of the above structures, except **T\_INFO**, contains at least one field of the **struct netbuf** type. The **netbuf** structure is described in the **tiuser.h** file. For each field of this type, the user may specify that the buffer for that field should be allocated as well. The *fields* parameter specifies this option, where the parameter is the bitwise-OR of any of the following:

**T\_ADDR**  
     The *addr* field of the **t\_bind**, **t\_call**, **t\_unitdata**, or **t\_uderr** structure.

**T\_OPT**   The *opt* field of the **t\_optmgmt**, **t\_call**, **t\_unitdata**, or **t\_uderr** structure.

**T\_UDATA**  
     The *udata* field of the **t\_call**, **t\_discon**, or **t\_unitdata** structure.

*fields*   **T\_ALL** All relevant fields of the given structure.

Specifies whether the buffer should be allocated for each field type. For each field specified in the *fields* parameter, the **t\_alloc** subroutine allocates memory for the buffer associated with the field, initializes the *len* field to zero, and initializes the *buf* pointer and the *maxlen* field accordingly. The length of the buffer allocated is based on the same size information returned to the user from the **t\_open** and **t\_getinfo** subroutines. Thus, the *fd* parameter must refer to the transport end point through which the newly allocated structure will be passed, so that the appropriate size information can be accessed. If the size value associated with any specified field is -1 or -2, the **t\_alloc** subroutine will be unable to determine the size of the buffer to allocate; it then fails, setting the **t\_errno** variable to **TSYSERR** and the **errno** global variable to **EINVAL**. For any field not specified in the *fields* parameter, the *buf* field is set to null and the *maxlen* field is set to 0.

## Return Values

On successful completion, the **t\_alloc** subroutine returns a pointer to the newly allocated structure. Otherwise, it returns a null pointer.

## Error Codes

On failure, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.

Value	Description
<b>TNOSTRUCTYPE</b>	Unsupported structure type requested. This can include a request for a structure type which is inconsistent with the transport provider type specified, for example, connection-oriented or connectionless.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_free** subroutine, **t\_getinfo** subroutine, **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_bind Subroutine for Transport Layer Interface

### Purpose

Binds an address to a transport end point.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_bind(fd, req, ret)
int fd;
struct t_bind * req;
struct t_bind * ret;
```

### Description

The **t\_bind** subroutine associates a protocol address with the transport end point specified by the *fd* parameter and activates that transport end point. In connection mode, the transport provider may begin accepting or requesting connections on the transport end point. In connectionless mode, the transport user may send or receive data units through the transport end point.

### Parameters

*fd* Specifies the transport end point.  
*req* Specifies the address to be bound to the given transport end point.  
*ret* Specifies the maximum size of the address buffer.

The *req* and *ret* parameters point to a **t\_bind** structure containing the following members:

```
struct netbuf addr;
unsigned qlen;
```

The **netbuf** structure is described in the **tiuser.h** file. The *addr* field of the **t\_bind** structure specifies a protocol address and the *qlen* field is used to indicate the maximum number of outstanding connect indications.

The *req* parameter is used to request that the address represented by the **netbuf** structure be bound to the given transport end point. In the *req* parameter, the **netbuf** structure fields have the following meanings:

Field	Description
len	Specifies the number of bytes in the address.
buf	Points to the address buffer.
maxlen	Has no meaning for the <i>req</i> parameter.

On return, the *ret* parameter contains the address that the transport provider actually bound to the transport end point; this may be different from the address specified by the user in the *req* parameter. In the *ret* parameter, the **netbuf** structure fields have the following meanings:

Field	Description
maxlen	Specifies the maximum size of the address buffer.
buf	Points to the buffer where the address is to be placed. (On return, this field points to the bound address.)
len	Specifies the number of bytes in the bound address.

If the value of the *maxlen* field is not large enough to hold the returned address, an error will result.

If the requested address is not available or if no address is specified in the *req* parameter (that is, the *len* field of the *addr* field in the *req* parameter is 0) the transport provider assigns an appropriate address to be bound and returns that address in the *addr* field of the *ret* parameter. The user can compare the addresses in the *req* parameter to those in the *ret* parameter to determine whether the transport provider has bound the transport end point to a different address than that requested. If the transport provider could not allocate an address, the **t\_bind** subroutine fails and **t\_errno** is set to **TNOADDR**.

The *req* parameter may be null if the user does not wish to specify an address to be bound. Here, the value of the *qlen* field is assumed to be 0, and the transport provider must assign an address to the transport end point. Similarly, the *ret* parameter may be null if the user does not care which address was bound by the provider and is not interested in the negotiated value of the *qlen* field. It is valid to set the *req* and *ret* parameters to null for the same call, in which case the provider chooses the address to bind to the transport end point and does not return that information to the user.

The *qlen* field has meaning only when initializing a connection-mode service. It specifies the number of outstanding connect indications the transport provider should support for the given transport end point. An outstanding connect indication is one that has been passed to the transport user by the transport provider. A value of the *qlen* field greater than 0 is only meaningful when issued by a passive transport user that expects other users to call it. The value of the *qlen* field is negotiated by the transport provider and can be changed if the transport provider cannot support the specified number of outstanding connect indications. On return, the *qlen* field in the *ret* parameter contains the negotiated value.

This subroutine allows more than one transport end point to be bound to the same protocol address as long as the transport provider also supports this capability. However, it is not allowable to bind more than one protocol address to the same transport end point. If a user binds more than one transport end point to the same protocol address, only one end point can be used to listen for connect indications associated with that protocol address. In other words, only one **t\_bind** subroutine for a given protocol address may specify a value greater than 0 for the *qlen* field. In this way, the transport provider can identify which transport end point should be notified of an incoming connect indication. If a user attempts to bind a protocol address to a second transport end point having a *qlen* value greater than 0, the transport provider instead assigns another address to be bound to that end point. If a user accepts a connection on the transport end point that is being used as the listening end point, the bound protocol address is found to be busy for the duration of that connection. No other transport end points may be bound for listening while that initial listening end point is in the data-transfer phase. This prevents more than one transport end point bound to the same protocol address from accepting connect indications.

## Return Values

On successful completion, the **t\_connect** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to use the specified address.
<b>TADDRBUSY</b>	The requested address is in use.
<b>TBADADDR</b>	The specified protocol address was in an incorrect format or contained illegal information.
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBUFOVFLW</b>	The number of bytes allowed for an incoming argument is not sufficient to store the value of that argument. The provider's state changes to <b>T_IDLE</b> and the information to be returned in the <i>ret</i> parameter is discarded.
<b>TNOADDR</b>	The transport provider could not allocate an address.
<b>TOUTSTATE</b>	The function was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_unbind** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_close Subroutine for Transport Layer Interface

### Purpose

Closes a transport end point.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_close(fd)
int fd;
```

### Description

The **t\_close** subroutine informs the transport provider that the user is finished with the transport end point specified by the *fd* parameter and frees any local library resources associated with the end point. In addition, the **t\_close** subroutine closes the file associated with the transport end point.

The **t\_close** subroutine should be called from the **T\_UNBND** state (see the **t\_getstate** subroutine). However, this subroutine does not check state information, so it may be called from any state to close a transport end point. If this occurs, the local library resources associated with the end point are freed automatically. In addition, the **close** subroutine is issued for that file descriptor. The **close** subroutine is abortive if no other process has that file open, and will break any transport connection that may be associated with that end point.

## Parameter

*fd* Specifies the transport end point to be closed.

## Return Values

On successful completion, the **t\_connect** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Code

If unsuccessful, the **t\_errno** variable is set to the following:

Value	Description
TBADF	The specified file descriptor does not refer to a transport end point.

## Related Information

The **close** subroutine, **t\_getstate** subroutine, **t\_open** subroutine, **t\_unbind** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_connect Subroutine for Transport Layer Interface

### Purpose

Establishes a connection with another transport user.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_connect(fd, sndcall, rcvcall)
int fd;
struct t_call * sndcall;
struct t_call * rcvcall;
```

### Description

The **t\_connect** subroutine enables a transport user to request a connection to the specified destination transport user.

### Parameters

<i>fd</i>	Identifies the local transport end point where communication will be established.
<i>sndcall</i>	Specifies information needed by the transport provider to establish a connection.
<i>rcvcall</i>	Specifies information associated with the newly established connection.

The *sndcall* and *rcvcall* parameters point to a **t\_call** structure that contains the following members:

```

struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;

```

The **netbuf** structure is described in the **tiuser.h** file. In the *sndcall* parameter, the **addr** field specifies the protocol address of the destination transport user, the **opt** field presents any protocol-specific information that might be needed by the transport provider, the **udata** field points to optional user data that may be passed to the destination transport user during connection establishment, and the **sequence** field has no meaning for this function.

On return to the *rcvcall* parameter, the **addr** field returns the protocol address associated with the responding transport end point, the **opt** field presents any protocol-specific information associated with the connection, the **udata** field points to optional user data that may be returned by the destination transport user during connection establishment; and the **sequence** field has no meaning for this function.

The **opt** field implies no structure on the options that may be passed to the transport provider. The transport provider is free to specify the structure of any options passed to it. These options are specific to the underlying protocol of the transport provider. The user can choose not to negotiate protocol options by setting the **len** field of the **opt** field to 0. In this case, the provider may use default options.

The **udata** field enables the caller to pass user data to the destination transport user and receive user data from the destination user during connection establishment. However, the amount of user data must not exceed the limits supported by the transport provider as returned by the **t\_open** or **t\_getinfo** subroutine. If the **len** field of the **udata** field in the *sndcall* parameter is 0, no data is sent to the destination transport user.

On return, the **addr**, **opt**, and **udata** fields of the *rcvcall* parameter are updated to reflect values associated with the connection. Thus, the **maxlen** field of each parameter must be set before issuing this function to indicate the maximum size of the buffer for each. However, the *rcvcall* parameter may be null, in which case no information is given to the user on return from the **t\_connect** subroutine.

By default, the **t\_connect** subroutine executes in synchronous mode, and waits for the destination user's response before returning control to the local user. A successful return (that is, a return value of 0) indicates that the requested connection has been established. However, if the **O\_NDELAY** flag is set (with the **t\_open** subroutine or the **fcntl** command), the **t\_connect** subroutine executes in asynchronous mode. In this case, the call does not wait for the remote user's response, but returns control immediately to the local user and returns -1 with the **t\_errno** variable set to **TNODATA** to indicate that the connection has not yet been established. In this way, the function simply initiates the connection establishment procedure by sending a connect request to the destination transport user.

## Return Values

On successful completion, the **t\_connect** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to use the specified address or options.
<b>TBADADDR</b>	The specified protocol address was in an incorrect format or contained illegal information.
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBADOPT</b>	The specified protocol options were in an incorrect format or contained illegal information.

Value	Description
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming argument is not sufficient to store the value of that argument. If executed in synchronous mode, the provider's state, as seen by the user, changes to <b>T_DATAXFER</b> , and the connect indication information to be returned in the <i>rcvcall</i> parameter is discarded.
<b>TLOOK</b>	An asynchronous event has occurred on this transport end point and requires immediate attention.
<b>TNODATA</b>	The <b>O_NDELAY</b> or <b>O_NONBLOCK</b> flag was set, so the function successfully initiated the connection establishment procedure, but did not wait for a response from the remote user.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The function was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **fcntl** command.

The **t\_accept** subroutine, **t\_getinfo** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_rcvconnect** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_error Subroutine for Transport Layer Interface

### Purpose

Produces an error message.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
void t_error(errmsg)
char * errmsg;
extern int t_errno;
extern char *t_errno;
extern int t_nerr;
```

### Description

The **t\_error** subroutine produces a message on the standard error output that describes the last error encountered during a call to a transport function.

The **t\_error** subroutine prints the user-supplied error message, followed by a colon and the standard transport-function error message for the current value contained in the **t\_errno** variable.

### Parameter

**errmsg** Specifies a user-supplied error message that gives context to the error.

## External Variables

<b>t_errno</b>	Specifies which standard transport-function error message to print. If the value of the <b>t_errno</b> variable is <b>TSYSERR</b> , the <b>t_error</b> subroutine also prints the standard error message for the current value contained in the <b>errno</b> global variable.
<b>t_errr</b>	The <b>t_errno</b> variable is set when an error occurs and is not cleared on subsequent successful calls. Specifies the maximum index value for the <b>t_errlist</b> array. The <b>t_errlist</b> array is the array of message strings allowing user-message formatting. The <b>t_errno</b> variable can be used as an index into this array to retrieve the error message string (without a terminating new-line character).

## Examples

A **t\_connect** subroutine is unsuccessful on transport end point fd2 because a bad address was given, and the following call follows the failure:

```
t_error("t_connect failed on fd2")
```

The diagnostic message would print as:

```
t_connect failed on fd2: Incorrect transport address format
```

In this example, **t\_connect** failed on fd2 tells the user which function was unsuccessful on which transport end point, and **Incorrect transport address format** identifies the specific error that occurred.

## Related Information

List of Streams Programming References, STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_free Subroutine for Transport Layer Interface

### Purpose

Frees a library structure.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_free(ptr, struct_type)
char * ptr;
int struct_type;
```

### Description

The **t\_free** subroutine frees memory previously allocated by the **t\_alloc** subroutine. This subroutine frees memory for the specified structure and also frees memory for buffers referenced by the structure.

The **t\_free** subroutine checks the **addr**, **opt**, and **udata** fields of the given structure (as appropriate) and frees the buffers pointed to by the **buf** field of the **netbuf** structure. If the **buf** field is null, the **t\_free** subroutine does not attempt to free memory. After all buffers are freed, the **t\_free** subroutine frees the memory associated with the structure pointed to by the **ptr** parameter.

Undefined results will occur if the **ptr** parameter or any of the **buf** pointers points to a block of memory that was not previously allocated by the **t\_alloc** subroutine.

## Parameters

*ptr* Points to one of the seven structure types described for the **t\_alloc** subroutine.  
*struct\_type* Identifies the type of that structure. The type can be one of the following:

<b>Type</b>	Structure
<b>T_BIND</b>	<b>struct t_bind</b>
<b>T_CALL</b>	<b>struct t_call</b>
<b>T_OPTMGMT</b>	<b>struct t_optmgmt</b>
<b>T_DIS</b>	<b>struct t_discon</b>
<b>T_UNITDATA</b>	<b>struct t_unitdata</b>
<b>T_UDERROR</b>	<b>struct t_uderr</b>
<b>T_INFO</b>	<b>struct t_info</b>

Each of these structure types is used as a parameter to one or more transport subroutines.

## Return Values

On successful completion, the **t\_free** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to the following:

Value	Description
<b>TNOSTRUCTYPE</b>	Unsupported structure type requested.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_alloc** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_getinfo Subroutine for Transport Layer Interface

### Purpose

Gets protocol-specific service information.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_getinfo(fd, info)  
int fd;  
struct t_info * info;
```

## Description

The **t\_getinfo** subroutine returns the current characteristics of the underlying transport protocol associated with *fd* file descriptor. The **t\_info** structure is used to return the same information returned by the **t\_open** subroutine. This function enables a transport user to access this information during any phase of communication.

## Parameters

*fd*        Specifies the file descriptor.

*info* Points to a **t\_info** structure that contains the following members:

```
long addr;  
long options;  
long tsdu;  
long etsdu;  
long connect;  
long discon;  
long servtype;
```

The values of the fields have the following meanings:

- addr** A value greater than or equal to 0 indicates the maximum size of a transport protocol address; a value of -1 specifies that there is no limit on the address size; and a value of -2 specifies that the transport provider does not provide user access to transport protocol addresses.
- options** A value greater than or equal to 0 indicates the maximum number of bytes of protocol-specific options supported by the provider; a value of -1 specifies that there is no limit on the option size; and a value of -2 specifies that the transport provider does not support user-settable options.
- tsdu** A value greater than 0 specifies the maximum size of a transport service data unit (TSDU); a value of 0 specifies that the transport provider does not support the concept of TSdu, although it does support the sending of a data stream having no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of a TSdu; and a value of -2 specifies that the transfer of normal data is not supported by the transport provider.
- etsdu** A value greater than 0 specifies the maximum size of an expedited transport service data unit (ETSDU); a value of 0 specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream having no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of an ETSDU; and a value of -2 specifies that the transfer of expedited data is not supported by the transport provider.
- connect** A value greater than or equal to 0 specifies the maximum amount of data that may be associated with connection establishment functions; a value of -1 specifies that there is no limit on the amount of data sent during connection establishment; and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment functions.
- discon** A value greater than or equal to 0 specifies the maximum amount of data that may be associated with the **t\_snddis** and **t\_rcvdis** subroutines; a value of -1 specifies that there is no limit on the amount of data sent with these abortive release functions; and a value of -2 specifies that the transport provider does not allow data to be sent with the abortive release functions.
- servtype** This field specifies the service type supported by the transport provider.

If a transport user is concerned with protocol independence, the sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the **t\_alloc** subroutine may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any function. The value of each field can change as a result of option negotiation; the **t\_getinfo** subroutine enables a user to retrieve the current characteristics.

## Return Values

On successful completion, the **t\_getinfo** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

The *servtype* field of the *info* parameter may specify one of the following values on return:

Value	Description
<b>T_COTS</b>	The transport provider supports a connection-mode service, but does not support the optional orderly release facility.

Value	Description
<b>T_COTS_ORD</b>	The transport provider supports a connection-mode service with the optional orderly release facility.
<b>T_CLTS</b>	The transport provider supports a connectionless-mode service. For this service type, the <b>t_open</b> subroutine returns -2 for the values in the <b>etsdu</b> , <b>connect</b> , and <b>discon</b> fields.

## Error Codes

In unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvdis** subroutine and **t\_snddis** subroutine.

List of Streams Programming Reference and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_getstate Subroutine for Transport Layer Interface

### Purpose

Gets the current state.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_getstate(fd)
int fd;
```

### Description

The **t\_getstate** subroutine returns the current state of the provider associated with the transport end point specified by the *fd* parameter.

### Parameter

*fd* Specifies the transport end point.

### Return Codes

On successful completion, the **t\_getstate** subroutine returns the current state. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

If the provider is undergoing a state transition when the **t\_getstate** subroutine is called, the function will fail. The current state is one of the following.

Value	Description
<b>T_DATAXFER</b>	Data transfer.
<b>T_IDLE</b>	Idle.
<b>T_INCON</b>	Incoming connection pending.
<b>T_INREL</b>	Incoming orderly release (waiting to send an orderly release indication).
<b>T_OUTCON</b>	Outgoing connection pending.
<b>T_OUTREL</b>	Outgoing orderly release (waiting for an orderly release indication).
<b>T_UNBND</b>	Unbound.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TSTATECHNG</b>	The transport provider is undergoing a state change.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_listen Subroutine for Transport Layer Interface

### Purpose

Listens for a connect request.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>

int t_listen(fd, call)
int fd;
struct t_call * call;
```

### Description

The **t\_listen** subroutine listens for a connect request from a calling transport user.

**Note:** If a user issues a **t\_listen** subroutine call in synchronous mode on a transport end point that was not bound for listening (that is, the `qlen` field was 0 on the **t\_bind** subroutine), the call will never return because no connect indications will arrive on that endpoint.

### Parameters

*fd* Identifies the local transport endpoint where connect indications arrive.

*call* Contains information describing the connect indication.

The *call* parameter points to a **t\_call** structure that contains the following members:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;  
int sequence;
```

The **netbuf** structure contains the following fields:

**addr** Returns the protocol address of the calling transport user.

**opt** Returns protocol-specific parameters associated with the connect request.

**udata** Returns any user data sent by the caller on the connect request.

**sequence**

Uniquely identifies the returned connect indication. The value of sequence enables the user to listen for multiple connect indications before responding to any of them.

Since the **t\_listen** subroutine returns values for the *addr*, *opt*, and *udata* fields of the *call* parameter, the *maxlen* field of each must be set before issuing the **t\_listen** subroutine to indicate the maximum size of the buffer for each.

By default, the **t\_listen** subroutine executes in synchronous mode and waits for a connect indication to arrive before returning to the user. However, if the **O\_NDELAY** or **O\_NONBLOCK** flag is set (using the **t\_open** subroutine or the **fcntl** command), the **t\_listen** subroutine executes asynchronously, reducing to a poll for existing connect indications. If none are available, the **t\_listen** subroutine returns -1 and sets the **t\_errno** variable to **TNODATA**.

## Return Values

On successful completion, the **t\_listen** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBADQLEN</b>	The transport end point is not bound for listening. The <i>qlen</i> is zero.
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming argument is not sufficient to store the value of that argument. The provider's state, as seen by the user, changes to <b>T_INCON</b> , and the connect-indication information to be returned in the <i>call</i> parameter is discarded.
<b>TLOOK</b>	An asynchronous event has occurred on this transport end point and requires immediate attention.
<b>TNODATA</b>	The <b>O_NDELAY</b> or <b>O_NONBLOCK</b> flag was set, but no connect indications had been queued.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_accept** subroutine, **t\_alloc** subroutine, **t\_bind** subroutine, **t\_connect** subroutine, **t\_open** subroutine, **t\_rcvconnect** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_look Subroutine for Transport Layer Interface

### Purpose

Looks at the current event on a transport end point.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_look(fd)
int fd;
```

### Description

The **t\_look** subroutine returns the current event on the transport end point specified by the *fd* parameter. This subroutine enables a transport provider to notify a transport user of an asynchronous event when the user is issuing functions in synchronous mode. Certain events require immediate notification of the user and are indicated by a specific error, **TLOOK**, on the current or next subroutine executed.

This subroutine also enables a transport user to poll a transport end point periodically for asynchronous events.

### Parameter

*fd* Specifies the transport end point.

### Return Values

On successful completion, the **t\_look** subroutine returns a value that indicates which of the allowable events has occurred, or returns a value of 0 if no event exists. One of the following events is returned:

Event	Description
<b>T_CONNECT</b>	Indicates connect confirmation received.
<b>T_DATA</b>	Indicates normal data received.
<b>T_DISCONNECT</b>	Indicates disconnect received.
<b>T_ERROR</b>	Indicates fatal error.
<b>T_EXDATA</b>	Indicates expedited data received.
<b>T_LISTEN</b>	Indicates connection indication received.
<b>T_ORDREL</b>	Indicates orderly release.
<b>T_UDERR</b>	Indicates datagram error.

If the **t\_look** subroutine is unsuccessful, a value of -1 is returned, and the **t\_errno** variable is set to indicate the error.

### Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_open Subroutine for Transport Layer Interface

### Purpose

Establishes a transport end point.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_open(path, oflag, info)
char * path;
int oflag;
struct t_info * info;
```

### Description

The **t\_open** subroutine must be called as the first step in the initialization of a transport end point. This subroutine establishes a transport end point, first, by opening a UNIX system file that identifies a particular transport provider (that is, transport protocol) and then returning a file descriptor that identifies that end point. For example, opening the **/dev/dlpi/tr** file identifies an 802.5 data link provider.

### Parameters

<i>path</i>	Points to the path name of the file to open.
<i>oflag</i>	Specifies the open routine flags.

*info* Points to a **t\_info** structure.

The *info* parameter points to a **t\_info** structure that contains the following elements:

```
long addr;  
long options;  
long tsdu;  
long etsdu;  
long connect;  
long discon;  
long servtype;
```

The values of the elements have the following meanings:

**addr** A value greater than or equal to 0 indicates the maximum size of a transport protocol address; a value of -1 specifies that there is no limit on the address size; and a value of -2 specifies that the transport provider does not provide user access to transport protocol addresses.

**options** A value greater than or equal to 0 indicates the maximum number of bytes of protocol-specific options supported by the provider; a value of -1 specifies that there is no limit on the option size; and a value of -2 specifies that the transport provider does not support user-settable options.

**tsdu** A value greater than 0 specifies the maximum size of a transport service data unit (TSDU); a value of 0 specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a data stream having no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of a TSDU; and a value of -2 specifies that the transfer of normal data is not supported by the transport provider.

**etsdu** A value greater than 0 specifies the maximum size of an expedited transport service data unit (ETSDU); a value of 0 specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream having no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of an ETSDU; and a value of -2 specifies that the transfer of expedited data is not supported by the transport provider.

**connect** A value greater than or equal to 0 specifies the maximum amount of data that may be associated with connection establishment functions; a value of -1 specifies that there is no limit on the amount of data sent during connection establishment; and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment functions.

**discon** A value greater than or equal to 0 specifies the maximum amount of data that may be associated with the **t\_snddis** and **t\_rcvdis** functions; a value of -1 specifies that there is no limit on the amount of data sent with these abortive release functions; and a value of -2 specifies that the transport provider does not allow data to be sent with the abortive release functions.

**servtype** This field specifies the service type supported by the transport provider, as described in the Return Values section.

If a transport user is concerned with protocol independence, these sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the **t\_alloc** subroutine can be used to allocate these buffers. An error results if a transport user exceeds the allowed data size on any function.

## Return Values

On successful completion, the **t\_open** subroutine returns a valid file descriptor. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

The *servtype* field of the *info* parameter can specify one of the following values on return:

Value	Description
<b>T_COTS</b>	The transport provider supports a connection-mode service but does not support the optional orderly release facility.
<b>T_COTS_ORD</b>	The transport provider supports a connection-mode service with the optional orderly release facility.
<b>T_CLTS</b>	The transport provider supports a connectionless-mode service. For this service type, the <b>t_open</b> subroutine returns -2 for the values in the <i>etsdu</i> , <i>connect</i> , and <i>discon</i> fields.

A single transport end point can support only one of the above services at one time.

If the *info* parameter is set to null by the transport user, no protocol information is returned by the **t\_open** subroutine.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to the following:

Value	Description
<b>TSYSERR</b>	A system error has occurred during the startup of this function.

## Related Information

The **open** subroutine, **t\_close** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_optmgmt Subroutine for Transport Layer Interface

### Purpose

Manages options for a transport end point.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_optmgmt(fd, req, ret)
int fd;
struct t_optmgmt * req;
struct t_optmgmt * ret;
```

### Description

The **t\_optmgmt** subroutine enables a transport user to retrieve, verify, or negotiate protocol options with the transport provider.

### Parameters

*fd*      Identifies a bound transport end point.  
*req*     Requests a specific action of the provider.

*ret* Returns options and flag values to the user.

Both the *req* and *ret* parameters point to a **t\_optmgmt** structure containing the following members:

```
struct netbuf opt;  
long flags;
```

The *opt* field identifies protocol options, and the *flags* field specifies the action to take with those options.

The options are represented by a **netbuf** structure in a manner similar to the address in the **t\_bind** subroutine. The *req* parameter is used to send options to the provider. This **netbuf** structure contains the following fields:

Field	Description
<i>len</i>	Specifies the number of bytes in the options.
<i>buf</i>	Points to the options buffer.
<i>maxlen</i>	Has no meaning for the <i>req</i> parameter.

The *ret* parameter is used to return information to the user from the transport provider. On return, this **netbuf** structure contains the following fields:

Field	Description
<i>len</i>	Specifies the number of bytes of options returned.
<i>buf</i>	Points to the buffer where the options are to be placed.
<i>maxlen</i>	Specifies the maximum size of the options buffer. The <i>maxlen</i> field has no meaning for the <i>req</i> parameter, but must be set in the <i>ret</i> parameter to specify the maximum number of bytes the options buffer can hold. The actual structure and content of the options is imposed by the transport provider.

The *flags* field of the *req* parameter can specify one of the following actions:

Action	Description
<b>T_NEGOTIATE</b>	Enables the user to negotiate the values of the options specified in the <i>req</i> parameter with the transport provider. The provider evaluates the requested options and negotiates the values, returning the negotiated values through the <i>ret</i> parameter.
<b>T_CHECK</b>	Enables the user to verify if the options specified in the <i>req</i> parameter are supported by the transport provider. On return, the <i>flags</i> field of the <i>ret</i> parameter has either <b>T_SUCCESS</b> or <b>T_FAILURE</b> set to indicate to the user whether the options are supported or not. These flags are only meaningful for the <b>T_CHECK</b> request.
<b>T_DEFAULT</b>	Enables a user to retrieve the default options supported by the transport provider into the <i>opt</i> field of the <i>ret</i> parameter. In the <i>req</i> parameter, the <i>len</i> field of the <i>opt</i> field must be zero, and the <i>buf</i> field can be NULL.

If issued as part of the connectionless-mode service, the **t\_optmgmt** subroutine may become blocked due to flow control constraints. The subroutine does not complete until the transport provider has processed all previously sent data units.

## Return Values

On successful completion, the **t\_optmgmt** subroutine returns a value of 0. Otherwise, it returns a value of -1, and the **t\_errno** variable is set to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TACCES</b>	User does not have permission to negotiate the specified options.
<b>TBADF</b>	Specified file descriptor does not refer to a transport endpoint.
<b>TBADFLAG</b>	Unusable flag was specified.
<b>TBADOPT</b>	Specified protocol options were in an incorrect format or contained unusable information.
<b>TBUFOVFLW</b>	Number of bytes allowed for an incoming parameter is not sufficient to store the value of that parameter. Information to be returned in the <i>ret</i> parameter will be discarded.
<b>TOUTSTATE</b>	Function was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during operation of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcv Subroutine for Transport Layer Interface

### Purpose

Receives normal data or expedited data sent over a connection.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
int t_rcv(fd, buf, nbytes, flags)
int fd;
char * buf;
unsigned nbytes;
int * flags;
```

### Description

The **t\_rcv** subroutine receives either normal or expedited data. By default, the **t\_rcv** subroutine operates in synchronous mode and will wait for data to arrive if none is currently available. However, if the **O\_NDELAY** flag is set (using the **t\_open** subroutine or the **fcntl** command), the **t\_rcv** subroutine runs in asynchronous mode and will stop if no data is available.

On return from the call, if the **T\_MORE** flag is set in the *flags* parameter, this indicates that there is more data. This means that the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple **t\_rcv** subroutine calls. Each **t\_rcv** subroutine with the **T\_MORE** flag set indicates that another **t\_rcv** subroutine must follow immediately to get more data for the current TSDU. The end of the TSDU is identified by the return of a **t\_rcv** subroutine call with the **T\_MORE** flag not set. If the transport provider does not support the concept of a TSDU as indicated in the *info* parameter on return from a **t\_open** or **t\_getinfo** subroutine, the **T\_MORE** flag is not meaningful and should be ignored.

On return, the data returned is expedited data if the **T\_EXPEDITED** flag is set in the *flags* parameter. If the number of bytes of expedited data exceeds the value in the *nbytes* parameter, the **t\_rcv** subroutine will set the **T\_EXPEDITED** and **T\_MORE** flags on return from the initial call. Subsequent calls to retrieve the

remaining ETSDU not have the **T\_EXPEDITED** flag set on return. The end of the ETSDU is identified by the return of a **t\_rcv** subroutine call with the **T\_MORE** flag not set.

If expedited data arrives after part of a TSDU has been retrieved, receipt of the remainder of the TSDU will be suspended until the ETSDU has been processed. Only after the full ETSDU has been retrieved (the **T\_MORE** flag is not set) will the remainder of the TSDU be available to the user.

## Parameters

<i>fd</i>	Identifies the local transport end point through which data will arrive.
<i>buf</i>	Points to a receive buffer where user data will be placed.
<i>nbytes</i>	Specifies the size of the receiving buffer.
<i>flags</i>	Specifies optional flags.

## Return Values

On successful completion, the **t\_rcv** subroutine returns the number of bytes it received. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable may be set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TLOOK</b>	An asynchronous event has occurred on this transport end point and requires immediate attention.
<b>TNODATA</b>	The <b>O_NDELAY</b> flag was set, but no data is currently available from the transport provider.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during operation of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_look** subroutine, **t\_open** subroutine, **t\_snd** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcvconnect Subroutine for Transport Layer Interface

### Purpose

Receives the confirmation from a connect request.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>

int t_rcvconnect(fd, call)
int fd;
struct t_call * call;
```

## Description

The **t\_rcvconnect** subroutine enables a calling transport user to determine the status of a previously sent connect request and is used in conjunction with **t\_connect** to establish a connection in asynchronous mode. The connection will be established on successful completion of this function.

## Parameters

*fd*        Identifies the local transport end point where communication will be established.  
*call*      Contains information associated with the newly established connection.

The *call* parameter points to a **t\_call** structure that contains the following elements:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;  
int sequence;
```

The **netbuf** structure contains the following elements:

**addr**      Returns the protocol address associated with the responding transport end point.  
**opt**       Presents protocol-specific information associated with the connection.  
**udata**    Points to optional user data that may be returned by the destination transport user during connection establishment.  
**sequence**      Has no meaning for this function.

The *maxlen* field of each parameter must be set before issuing this function to indicate the maximum size of the buffer for each. However, the *call* parameter may be null, in which case no information is given to the user on return from the **t\_rcvconnect** subroutine. By default, the **t\_rcvconnect** subroutine runs in synchronous mode and waits for the connection to be established before returning. On return, the *addr*, *opt*, and *udata* fields reflect values associated with the connection.

If the **O\_NDELAY** flag is set (using the **t\_open** subroutine or **fcntl** command), the **t\_rcvconnect** subroutine runs in asynchronous mode and reduces to a poll for existing connect confirmations. If none are available, the **t\_rcvconnect** subroutine stops and returns immediately without waiting for the connection to be established. The **t\_rcvconnect** subroutine must be re-issued at a later time to complete the connection establishment phase and retrieve the information returned in the *call* parameter.

## Return Values

On successful completion, the **t\_rcvconnect** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable may be set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming parameter is not sufficient to store the value of that parameter and the connect information to be returned in the <i>call</i> parameter will be discarded. The state of the provider, as seen by the user, will be changed to <b>DATAXFER</b> .
<b>TLOOK</b>	An asynchronous event has occurred on this transport connection and requires immediate attention.
<b>TNODATA</b>	The <b>O_NDELAY</b> flag was set, but a connect confirmation has not yet arrived.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	This subroutine was issued in the wrong sequence.

Value	Description
TSYSERR	A system error has occurred during operation of this subroutine.

## Related Information

The **t\_accept** subroutine, **t\_alloc** subroutine, **t\_bind** subroutine, **t\_connect** subroutine, **t\_listen** subroutine, **t\_look** subroutine, **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcvdis Subroutine for Transport Layer Interface

### Purpose

Retrieves information from disconnect.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
t_rcvdis(fd, discon)
int fd;
struct t_discon * discon;
```

### Description

The **t\_rcvdis** subroutine is used to identify the cause of a disconnect, and to retrieve any user data sent with the disconnect.

### Parameters

*fd* Identifies the local transport end point where the connection existed.

*discon* Points to a **t\_discon** structure that contains the reason for the disconnect and contains any user data that was sent with the disconnect.

The **t\_discon** structure contains the following members:

```
struct netbuf udata;  
int reason;  
int sequence;
```

These fields are defined as follows:

**reason** Specifies the reason for the disconnect through a protocol-dependent reason code.

**udata** Identifies any user data that was sent with the disconnect.

**sequence**

Identifies an outstanding connect indication with which the disconnect is associated. The sequence field is only meaningful when the **t\_rcvdis** subroutine is issued by a passive transport user that has called one or more **t\_listen** subroutines and is processing the resulting connect indications. If a disconnect indication occurs, the sequence field can be used to identify which of the outstanding connect indications is associated with the disconnect.

If a user does not care if there is incoming data and does not need to know the value of the reason or sequence fields, the *discon* parameter may be null and any user data associated with the disconnect will be discarded. However, if a user has retrieved more than one outstanding connect indication (using the **t\_listen** subroutine) and the *discon* parameter is null, the user will be unable to identify with which connect indication the disconnect is associated.

## Return Values

On successful completion, the **t\_rcvdis** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable may be set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBUFOVFLW</b>	The number of bytes allocated for incoming data is not sufficient to store the data. (The state of the provider, as seen by the user, will change to <b>T_IDLE</b> , and the disconnect indication information to be returned in the <i>discon</i> parameter will be discarded.)
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNODIS</b>	No disconnect indication currently exists on the specified transport end point.
<b>TNOTSUPPORT</b>	This function is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	This subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_alloc** subroutine, **t\_connect** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_snddis** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcvrel Subroutine for Transport Layer Interface

### Purpose

Acknowledges receipt of an orderly release indication.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
t_rcvrel(fd)
int fd;
```

### Description

The **t\_rcvrel** subroutine is used to acknowledge receipt of an orderly release indication. After receipt of this indication, the user may not attempt to receive more data because such an attempt will block forever. However, the user may continue to send data over the connection if the **t\_sndrel** subroutine has not been issued by the user. The subroutine is an optional service of the transport provider, and is only supported if the transport provider returned service type **T\_COTS\_ORD** on the **t\_open** or **t\_getinfo** subroutine.

### Parameter

*fd*      Identifies the local transport end point where the connection exists.

### Return Values

On successful completion, the **t\_rcvrel** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

### Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TLOOK</b>	An asynchronous event has occurred on this transport end point and requires immediate attention.
<b>TNOREL</b>	No orderly release indication currently exists on the specified transport end point.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	This subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

### Related Information

The **t\_getinfo** subroutine, **t\_look** subroutine, **t\_open** subroutine, **t\_sndrel** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcvudata Subroutine for Transport Layer Interface

### Purpose

Receives a data unit.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_rcvudata(fd, unitdata, flags)
int fd;
struct t_unitdata * unitdata;
int * flags;
```

### Description

The **t\_rcvudata** subroutine is used in connectionless mode to receive a data unit from another transport user.

### Parameters

*fd* Identifies the local transport end point through which data will be received.

*unitdata* Holds information associated with the received data unit.

The *unitdata* parameter points to a **t\_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

On return from this call:

**addr** Specifies the protocol address of the sending user.

**opt** Identifies protocol-specific options that were associated with this data unit.

**udata** Specifies the user data that was received.

**Note:** The *maxlen* field of the *addr*, *opt*, and *udata* fields must be set before issuing this function to indicate the maximum size of the buffer for each.

*flags* Indicates that the complete data unit was not received.

By default, the **t\_rcvudata** subroutine operates in synchronous mode and will wait for a data unit to arrive if none is currently available. However, if the **O\_NDELAY** or **O\_NONBLOCK** flag is set (using the **t\_open** subroutine or **fcntl** command), the **t\_rcvudata** subroutine will run in asynchronous mode and will stop if no data units are available.

If the buffer defined in the *udata* field of *unitdata* is not large enough to hold the current data unit, the buffer will be filled and the **T\_MORE** flag will be set in *flags* on return to indicate that another **t\_rcvudata** subroutine should be issued to retrieve the rest of the data unit. Subsequent **t\_rcvudata** subroutine calls will return 0 for the length of the address and options until the full data unit has been received.

### Return Values

On successful completion, the **t\_rcvudata** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBUFOVFLW</b>	The number of bytes allocated for the incoming protocol address of options is not sufficient to store the information. (The unit data information to be returned in the <i>unitdata</i> parameter will be discarded.)
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNODATA</b>	The <b>O_DELAY</b> or <b>O_NONBLOCK</b> flag was set, but no data units are currently available from the transport provider.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during operation of this subroutine.

## Related Information

The **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvuderr** subroutine, **t\_sndudata** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_rcvuderr Subroutine for Transport Layer Interface

### Purpose

Receives a unit data error indication.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_rcvuderr(fd, uderr)
int fd;
struct t_uderr * uderr;
```

### Description

The **t\_rcvuderr** subroutine is used in connectionless mode to receive information concerning an error on a previously sent data unit, and should only be issued following a unit data error indication. It informs the transport user that a data unit with a specific destination address and protocol options produced an error.

### Parameters

**fd** Identifies the local transport endpoint through which the error report will be received.

**uderr** Points to a **t\_uderr** structure containing the following members:

```
struct netbuf addr;  
struct netbuf opt;  
long error;
```

The **maxlen** field of the **addr** and **opt** fields must be set before issuing this function to indicate the maximum size of the buffer for each.

On return from this call, the **t\_uderr** structure contains:

**addr** Specifies the destination protocol address of the erroneous data unit.  
**opt** Identifies protocol-specific options that were associated with the data unit.  
**error** Specifies a protocol-dependent error code.

If the user decides not to identify the data unit that produced an error, the **uderr** parameter can be set to null and the **t\_rcvuderr** subroutine will clear the error indication without reporting any information to the user.

## Return Values

On successful completion, the **t\_rcvuderr** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TNOUDERR</b>	No unit data error indication currently exists on the specified transport end point.
<b>TBUFOVFLW</b>	The number of bytes allocated for the incoming protocol address or options is not sufficient to store the information. (The unit data error information to be returned in the <b>uderr</b> parameter will be discarded.)
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_look** subroutine, **t\_rcvudata** subroutine, **t\_sndudata** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_snd Subroutine for Transport Layer Interface

### Purpose

Sends data or expedited data over a connection.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_snd(fd, buf, nbytes, flags)
int fd;
char * buf;
unsigned nbytes;
int flags;
```

## Description

The **t\_snd** subroutine is used to send either normal or expedited data.

By default, the **t\_snd** subroutine operates in synchronous mode and may wait if flow-control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if the **O\_NDELAY** or **O\_NONBLOCK** flag is set (using the **t\_open** subroutine or the **fcntl** command), the **t\_snd** subroutine runs in asynchronous mode and stops immediately if there are flow-control restrictions.

Even when there are no flow-control restrictions, the **t\_snd** subroutine will wait if STREAMS internal resources are not available, regardless of the state of the **O\_NDELAY** or **O\_NONBLOCK** flag.

On successful completion, the **t\_snd** subroutine returns the number of bytes accepted by the transport provider. Normally this equals the number of bytes specified in the *nbytes* parameter. However, if the **O\_NDELAY** or **O\_NONBLOCK** flag is set, it is possible that only part of the data will be accepted by the transport provider. In this case, the **t\_snd** subroutine sets the **T\_MORE** flag for the data that was sent and returns a value less than the value of the *nbytes* parameter. If the value of the *nbytes* parameter is 0, no data is passed to the provider and the **t\_snd** subroutine returns a value of 0.

## Parameters

<i>fd</i>	Identifies the local transport end point through which data is sent.
<i>buf</i>	Points to the user data.
<i>nbytes</i>	Specifies the number of bytes of user data to be sent.
<i>flags</i>	Specifies any optional flags.

If the **T\_EXPEDITED** flag is set in the *flags* parameter, the data is sent as expedited data and is subject to the interpretations of the transport provider.

If the **T\_MORE** flag is set in the *flags* parameter, or as described above, an indication is sent to the transport provider that the transport service data unit (TSDU) or expedited transport service data unit (ETSDU) is being sent through multiple **t\_snd** subroutine calls. Each **t\_snd** subroutine with the **T\_MORE** flag set indicates that another **t\_snd** subroutine will follow with more data for the current TSDU. The end of the TSDU or ETSDU is identified by a **t\_snd** subroutine call with the **T\_MORE** flag not set. Use of the **T\_MORE** flag enables a user to break up large logical data units without losing the boundaries of those units at the other end of the connection. The flag implies nothing about how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU as indicated in the *info* parameter on return from the **t\_open** or **t\_getinfo** subroutine, the **T\_MORE** flag is not meaningful and should be ignored.

The size of each TSDU or ETSDU must not exceed the limits of the transport provider as returned by the **t\_open** or **t\_getinfo** subroutine. If the size is exceeded, a **TSYSERR** error with system error **EPROTO** occurs. However, the **t\_snd** subroutine may not fail because **EPROTO** errors may not be reported immediately. In this case, a subsequent call that accesses the transport endpoint fails with the associated **TSYSERR** error.

If the call to the **t\_snd** subroutine is issued from the **T\_IDLE** state, the provider may silently discard the data. If the call to the **t\_snd** subroutine is issued from any state other than **T\_DATAXFER**, **T\_INREL**, or **T\_IDLE**, the provider generates a **TSYSERR** error with system error **EPROTO** (which can be reported in the manner described above).

## Return Values

On successful completion, the **t\_snd** subroutine returns the number of bytes accepted by the transport provider. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The specified file descriptor does not refer to a transport end point.
<b>TBADFLAG</b>	The value specified in the <i>flags</i> parameter is invalid.
<b>TFLOW</b>	The <b>O_NDELAY</b> or <b>O_NONBLOCK</b> flag was set, but the flow-control mechanism prevented the transport provider from accepting data at this time.
<b>TLOOK</b>	An asynchronous event has occurred on the transport end point reference by the <i>fd</i> parameter and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has been detected during execution of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_getstate** subroutine, **t\_open** subroutine, **t\_rcv** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_snddis Subroutine for Transport Layer Interface

### Purpose

Sends a user-initiated disconnect request.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>

int t_snddis(fd, call)
int fd;
struct t_call * call;
```

### Description

The **t\_snddis** subroutine is used to initiate an abortive release on an already established connection or to reject a connect request.

### Parameters

*fd*        Identifies the local transport endpoint of the connection.

*call* Specifies information associated with the abortive release.

The *call* parameter points to a **t\_call** structure containing the following fields:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;  
int sequence;
```

The values in the *call* parameter have different semantics, depending on the context of the call to the **t\_snddis** subroutine. When rejecting a connect request, the *call* parameter must not be null and must contain a valid value in the sequence field to uniquely identify the rejected connect indication to the transport provider. The *addr* and *opt* fields of the *call* parameter are ignored. In all other cases, the *call* parameter need only be used when data is being sent with the disconnect request. The *addr*, *opt*, and sequence fields of the **t\_call** structure are ignored. If the user does not wish to send data to the remote user, the value of the *call* parameter can be null.

The *udata* field specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider as returned by the **t\_open** or **t\_getinfo** subroutine. If the *len* field of the *udata* field is 0, no data will be sent to the remote user.

## Return Values

On successful completion, the **t\_snddis** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TOUTSTATE</b>	The error code is returned when the communication endpoint referenced by the file descriptor is not in a state in which a call to this function is valid.
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider. The transport provider's outgoing queue will be flushed, so data might be lost.
<b>TBADSEQ</b>	An incorrect sequence number was specified, or a null call structure was specified when rejecting a connect request. The transport provider's outgoing queue will be flushed, so data might be lost.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_connect** subroutine, **t\_getinfo** subroutine, **t\_listen** subroutine, **t\_look** subroutine, **t\_open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_sndrel Subroutine for Transport Layer Interface

### Purpose

Initiates an orderly release of a transport connection.

## Library

Transport Layer Interface Library (**libtli.a**)

## Syntax

```
#include <tiuser.h>
```

```
int t_sndrel(fd)
int fd;
```

## Description

The **t\_sndrel** subroutine is used to initiate an orderly release of a transport connection and indicates to the transport provider that the transport user has no more data to send.

After issuing a **t\_sndrel** subroutine call, the user cannot send any more data over the connection. However, a user can continue to receive data if an orderly release indication has been received.

The **t\_sndrel** subroutine is an optional service of the transport provider and is only supported if the transport provider returned service type **T\_COTS\_ORD** in the **t\_open** or **t\_getinfo** subroutine.

## Parameter

*fd* Identifies the local transport endpoint where the connection exists.

## Return Values

On successful completion, the **t\_sndrel** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TFLOW</b>	The <b>O_NDELAY</b> or <b>O_NONBLOCK</b> flag was set, but the flow-control mechanism prevented the transport provider from accepting the function at this time.
<b>TLOOK</b>	An asynchronous event has occurred on the transport end point reference by the <i>fd</i> parameter and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine, **t\_rcvrel** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_sndudata Subroutine for Transport Layer Interface

### Purpose

Sends a data unit to another transport user.

## Library

Transport Layer Interface Library (**libtli.a**)

## Syntax

```
#include <tiuser.h>
```

```
int t_sndudata(fd, unitdata)
int fd;
struct t_unitdata * unitdata;
```

## Description

The **t\_sndudata** subroutine is used in connectionless mode to send a data unit to another transport user.

By default, the **t\_sndudata** subroutine operates in synchronous mode and may wait if flow-control restrictions prevent the data from being accepted by the local transport provider at the time the call is made. However, if the **O\_NDELAY** or **O\_NONBLOCK** flag is set (using the **t\_opensubroutine** or the **fcntl** command), the **t\_sndudata** subroutine runs in asynchronous mode and fails under such conditions.

## Parameters

*fd* Identifies the local transport endpoint through which data is sent.  
*unitdata* Points to a **t\_unitdata** structure containing the following elements:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

The elements are defined as follows:

**addr** Specifies the protocol address of the destination user.  
**opt** Identifies protocol-specific options that the user wants associated with this request.  
**udata** Specifies the user data to be sent. The user can choose not to specify what protocol options are associated with the transfer by setting the **len** field of the **opt** field to 0. In this case, the provider can use default options.

If the **len** field of the **udata** field is 0, no data unit is passed to the transport provider; the **t\_sndudata** subroutine does not send zero-length data units.

If the **t\_sndudata** subroutine is issued from an invalid state, or if the amount of data specified in the **udata** field exceeds the TSDU size as returned by the **t\_open** or **t\_getinfo** subroutine, the provider generates an **EPROTO** protocol error.

## Return Values

On successful completion, the **t\_sndudata** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TFLOW</b>	The <b>O_NDELAY</b> or <b>O_NONBLOCK</b> flag was set, but the flow-control mechanism prevented the transport provider from accepting data at this time.

Value	Description
<b>TLOOK</b>	An asynchronous event has occurred on this transport end point and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	This subroutine was issued in the wrong sequence.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvudata** subroutine, **t\_rcvuderr** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_sync Subroutine for Transport Layer Interface

### Purpose

Synchronizes transport library.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_sync(fd)
int fd;
```

### Description

The **t\_sync** subroutine synchronizes the data structures managed by the transport library with information from the underlying transport provider. In doing so, this subroutine can convert a raw file descriptor (obtained using the **open** or **dup** subroutine, or as a result of a **fork** operation and an **exec** operation) to an initialized transport endpoint, assuming that the file descriptor referenced a transport provider. This subroutine also allows two cooperating processes to synchronize their interaction with a transport provider.

For example, a process creates a new process with the **fork** subroutine and issues an **exec** subroutine call. The new process must issue a **t\_sync** subroutine call to build the private library data structure associated with a transport endpoint and to synchronize the data structure with the relevant provider information.

**Note:** The transport provider treats all users of a transport endpoint as a single user. If multiple processes are using the same endpoint, they should coordinate their activities so as not to violate the state of the provider. The **t\_sync** subroutine returns the current state of the provider to the user, thereby enabling the user to verify the state before taking further action. This coordination is only valid among cooperating processes; a process or an incoming event may change the provider's state *after* a **t\_sync** subroutine call is issued.

If the provider is undergoing a state transition when the **t\_sync** subroutine is called, the subroutine will be unsuccessful.

## Parameters

*fd* Specifies the transport end point.

## Return Values

On successful completion, the **t\_sync** subroutine returns the state of the transport provider. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error. The state returned can be one of the following:

Value	Description
<b>T_UNBIND</b>	Unbound
<b>T_IDLE</b>	Idle
<b>T_OUTCON</b>	Outgoing connection pending
<b>T_INCON</b>	Incoming connection pending
<b>T_DATAXFER</b>	Data transfer
<b>T_OUTREL</b>	Outgoing orderly release (waiting for an orderly release indication)
<b>T_INREL</b>	Incoming orderly release (waiting for an orderly release request)

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor is a valid open file descriptor, but does not refer to a transport endpoint.
<b>TSTATECHNG</b>	The transport provider is undergoing a state change.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **dup** subroutine, **exec** subroutine, **fork** subroutine, **open** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_unbind Subroutine for Transport Layer Interface

### Purpose

Disables a transport endpoint.

### Library

Transport Layer Interface Library (**libtli.a**)

### Syntax

```
#include <tiuser.h>
```

```
int t_unbind(fd)
int fd;
```

## Description

The **t\_unbind** subroutine disables a transport endpoint, which was previously bound by the **t\_bind** subroutine. On completion of this call, no further data or events destined for this transport endpoint are accepted by the transport provider.

## Parameter

*fd* Specifies the transport endpoint.

## Return Values

On successful completion, the **t\_unbind** subroutine returns a value of 0. Otherwise, it returns a value of -1 and sets the **t\_errno** variable to indicate the error.

## Error Codes

If unsuccessful, the **t\_errno** variable is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TOUTSTATE</b>	The function was issued in the wrong sequence.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **t\_bind** subroutine.

List of Streams Programming References and STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

---

## testb Utility

### Purpose

Checks for an available buffer.

### Syntax

```
int
testb(size, pri)
register size;
uint pri;
```

### Description

The **testb** utility checks for the availability of a message buffer of the size specified in the *size* parameter without actually retrieving the buffer. A successful return value from the **testb** utility does not guarantee that a subsequent call to the **allocb** utility will succeed; for example, when an interrupt routine takes the buffers.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*size* Specifies the buffer size.

*pri* Specifies the relative importance of the allocated blocks to the module. The possible values are:

- **BPRI\_LO**
- **BPRI\_MED**
- **BPRI\_HI**

The *pri* parameter is currently unused and is maintained only for compatibility with applications developed prior to UNIX System V Release 4.0.

## Return Values

If the buffer is available, the **testb** utility returns a value of 1. Otherwise, it returns a value of 0.

## Related Information

The **allocb** utility.

List of Streams Programming References and Understanding STREAMS Flow Control in *AIX Version 6.1 Communications Programming Concepts*.

---

## timeout Utility

### Purpose

Schedules a function to be called after a specified interval.

### Syntax

```
int
timeout(func, arg, ticks)
int (*func)();
caddr_t arg;
long ticks;
```

### Description

The **timeout** utility schedules the function pointed to by the *func* parameter to be called with the *arg* parameter after the number of timer ticks specified by the *ticks* parameter. Multiple pending calls to the **timeout** utility with the same *func* and *arg* parameters are allowed. The function called by the **timeout** utility must adhere to the same restrictions as a driver interrupt handler. It must not sleep.

On multiprocessor systems, the function called by the **timeout** utility should be interrupt-safe. Otherwise, the **STR\_QSAFETY** flag must be set when installing the module or driver with the **str\_install** utility.

This utility is part of STREAMS Kernel Extension.

**Note:** This utility must not be confused with the kernel service of the same name in the **libsys.a** library. STREAMS modules and drivers inherently use this version, not the **libsys.a** library version. No special action is required to use this version in the STREAMS environment.

### Parameters

*func* Indicates the function to be called. The function is declared as follows:

```
void (*func)(arg)
void *arg;
```

*arg* Indicates the parameter to supply to the function specified by the *func* parameter.

*ticks* Specifies the number of timer ticks that must occur before the function specified by the *func* parameter is called. Many timer ticks can occur every second.

## Return Values

The **timeout** utility returns an integer that identifies the request. This value may be used to withdraw the time-out request by using the **untimeout** utility. If the timeout table is full, the **timeout** utility returns a value of 0 and the request is not registered.

## Execution Environment

The **timeout** utility may be called from either the process or interrupt environment.

## Related Information

The **untimeout** utility.

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts*.

Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

Understanding STREAMS Synchronization in *AIX Version 6.1 Communications Programming Concepts* .

---

## timod Module

### Purpose

Converts a set of **streamio** operations into STREAMS messages.

### Description

The **timod** module is a STREAMS module for use with the Transport Interface (TI) functions of the Network Services Library. The **timod** module converts a set of **streamio** operations into STREAMS messages that may be consumed by a transport protocol provider that supports the Transport Interface. This allows a user to initiate certain TI functions as atomic operations.

The **timod** module must only be pushed (see "Pushable Modules" in *AIX Version 6.1 Communications Programming Concepts*) onto a stream terminated by a transport protocol provider that supports the TI.

All STREAMS messages, with the exception of the message types generated from the **streamio** operations described below as values for the *cmd* field, will be transparently passed to the neighboring STREAMS module or driver. The messages generated from the following **streamio** operations are recognized and processed by the **timod** module.

This module is part of STREAMS Kernel Extensions.

## Fields

The fields are described as follows:

Field	Description
-------	-------------

cmd	Specifies the command to be carried out. The possible values for this field are:
-----	--

<b>TI_BIND</b>	
----------------	--

	Binds an address to the underlying transport protocol provider. The message issued to the <b>TI_BIND</b> operation is equivalent to the TI message type <b>T_BIND_REQ</b> , and the message returned by the successful completion of the operation is equivalent to the TI message type <b>T_BIND_ACK</b> .
--	---

<b>TI_UNBIND</b>	
------------------	--

	Unbinds an address from the underlying transport protocol provider. The message issued to the <b>TI_UNBIND</b> operation is equivalent to the TI message type <b>T_UNBIND_REQ</b> , and the message returned by the successful completion of the operation is equivalent to the TI message type <b>T_OK_ACK</b> .
--	---

<b>TI_GETINFO</b>	
-------------------	--

	Gets the TI protocol-specific information from the transport protocol provider. The message issued to the <b>TI_GETINFO</b> operation is equivalent to the TI message type <b>T_INFO_REQ</b> , and the message returned by the successful completion of the operation is equivalent to the TI message type <b>T_INFO_ACK</b> .
--	--

<b>TI_OPTMGMT</b>	
-------------------	--

	Gets, sets, or negotiates protocol-specific options with the transport protocol provider. The message issued to the <b>TI_OPTMGMT</b> ioctl operation is equivalent to the TI message type <b>T_OPTMGMT_REQ</b> , and the message returned by the successful completion of the ioctl operation is equivalent to the TI message type <b>T_OPTMGMT_ACK</b> .
--	--

len	(On issuance) Specifies the size of the appropriate TI message to be sent to the transport provider.
-----	--

	(On return) Specifies the size of the appropriate TI message from the transport provider in response to the issued TI message.
--	--

dp	Specifies a pointer to a buffer large enough to hold the contents of the appropriate TI messages. The TI message types are defined in the <b>sys/tihdr.h</b> file.
----	--

## Examples

The following is an example of how to use the **timod** module:

```
#include <sys/stropts.h>
-
-
struct strioctl strioctl;
strucu t_info info;
-
-
strioctl.ic_cmd = TI_GETINFO;
strioctl.ic_timeout = INFTIM;
strioctl.ic_len = sizeof (info);
strioctl.ic_dp = (char *)&info;
ioctl(filides, I_STR, &strioctl);
```

## Related Information

The **tirdwr** module.

The **streamio** operations.

Benefits and Features of STREAMS, Building STREAMS, Pushable Modules, Understanding STREAMS Drivers and Modules, Understanding STREAMS Messages, Using STREAMS in *AIX Version 6.1 Communications Programming Concepts*.

---

## tirdwr Module

### Purpose

Supports the Transport Interface functions of the Network Services library.

### Description

The **tirdwr** module is a STREAMS module that provides an alternate interface to a transport provider that supports the Transport Interface (TI) functions of the Network Services library. This alternate interface allows a user to communicate with the transport protocol provider by using the **read** and **write** subroutines. The **putmsg** and **getmsg** system calls can also be used. However, the **putmsg** and **getmsg** system calls can only transfer data messages between user and stream.

The **tirdwr** module must only be pushed (see the **I\_PUSH** operation) onto a stream terminated by a transport protocol provider that supports the TI. After the **tirdwr** module has been pushed onto a stream, none of the TI functions can be used. Subsequent calls to TI functions will cause an error on the stream. Once the error is detected, subsequent system calls on the stream will return an error with the **errno** global variable set to **EPROTO**.

The following list describes actions taken by the **tirdwr** module when it is pushed or popped or when data passes through it:

Action	Description
<b>push</b>	Checks any existing data to ensure that only regular data messages are present. It ignores any messages on the stream that relate to process management. If any other messages are present, the <b>I_PUSH</b> operation returns an error and sets the <b>errno</b> global variable to <b>EPROTO</b> .
<b>write</b>	Takes the following actions on data that originated from a <b>write</b> subroutine:  <b>Messages with no control portions</b> Passes the message on downstream.  <b>Zero length data messages</b> Frees the message and does not pass downstream.  <b>Messages with control portions</b> Generates an error, fails any further system calls, and sets the <b>errno</b> global variable to <b>EPROTO</b> .
<b>read</b>	Takes the following actions on data that originated from the transport protocol provider:  <b>Messages with no control portions</b> Passes the message on upstream.  <b>Zero length data messages</b> Frees the message and does not pass upstream.  Messages with control portions will produce the following actions: <ul style="list-style-type: none"><li>• Messages that represent expedited data generate an error. All further calls associated with the stream fail with the <b>errno</b> global variable set to <b>EPROTO</b>.</li><li>• Any data messages with control portions have the control portions removed from the message prior to passing the message to the upstream neighbor.</li><li>• Messages that represent an orderly release indication from the transport provider generate a zero length data message, indicating the end of file, which is sent to the reader of the stream. The orderly release message itself is freed by the module.</li><li>• Messages that represent an abortive disconnect indication from the transport provider cause all further <b>write</b> and <b>putmsg</b> calls to fail with the <b>errno</b> global variable set to <b>ENXIO</b>. All further <b>read</b> and <b>getmsg</b> calls return zero length data (indicating end of file) once all previous data has been read.</li><li>• With the exception of the above rules, all other messages with control portions generate an error, and all further system calls associated with the stream fail with the <b>errno</b> global variable set to <b>EPROTO</b>.</li></ul>

Action	Description
<b>pop</b>	Sends an orderly release request to the remote side of the transport connection if an orderly release indication has been previously received.

## Related Information

The **timod** module.

The **streamio** operations.

The **read** subroutine, **write** subroutine.

The **getmsg** system call, **putmsg** system call.

Benefits and Features of STREAMS, Building STREAMS, Pushable Modules, STREAMS Overview, Understanding STREAMS Drivers and Modules, Understanding STREAMS Messages, Using STREAMS in *AIX Version 6.1 Communications Programming Concepts*.

---

## unbufcall Utility

### Purpose

Cancels a **bufcall** request.

### Syntax

```
void unbufcall(id)
register int id;
```

### Description

The **unbufcall** utility cancels a **bufcall** request.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*id* Identifies an event in the **bufcall** request.

## Related Information

The **bufcall** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## unlinkb Utility

### Purpose

Removes a message block from the head of a message.

## Syntax

```
mb1k_t *  
unlinkb(bp)  
register mb1k_t * bp;
```

## Description

The **unlinkb** utility removes the first message block pointed to by the *bp* parameter and returns a pointer to the head of the resulting message. The **unlinkb** utility returns a null pointer if there are no more message blocks in the message.

This utility is part of STREAMS Kernel Extensions.

## Parameters

*bp* Specifies which message block to unlink.

## Related Information

The **linkb** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## untimeout Utility

### Purpose

Cancels a pending timeout request.

### Syntax

```
int  
untimeout(id)  
int id;
```

### Description

The **untimeout** utility cancels the specific request made with the **timeout** utility.

This utility is part of STREAMS Kernel Extensions.

**Note:** This utility must not be confused with the kernel service of the same name in the **libsys.a** library. STREAMS modules and drivers inherently use this version, not the **libsys.a** library version. No special action is required to use this version in the STREAMS environment.

### Parameters

*id* Specifies the identifier returned from the corresponding timeout request.

### Execution Environment

The **untimeout** utility can be called from either the process or interrupt environment.

## Related Information

The **timeout** utility.

List of Streams Programming References and Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

---

## unweldq Utility

### Purpose

Removes a previously established weld connection between STREAMS queues.

### Syntax

```
#include <sys/stream.h>
```

```
int unweldq ( q1, q2, q3, q4, func, arg, protect_q)
queue_t *q1;
queue_t *q2;
queue_t *q3;
queue_t *q4;
weld_fcn_t func;
weld_arg_t arg;
queue_t *protect_q;
```

### Description

The **unweldq** utility removes a weld connection previously established with the **weld** utility between two STREAMS queues (*q1* and *q2*). The **unweldq** utility can be used to unweld two pairs of queues in one call (*q1* and *q2*, *q3* and *q4*).

The unwelding operation is performed by changing the first queue's **q\_next** pointer so that it does not point to any queue. The **unweldq** utility does not actually perform the operation. Instead, it creates an unwelding request which STREAMS performs asynchronously. STREAMS acquires the appropriate synchronization queues before performing the operation.

Callers that need to know when the unwelding operation has actually taken place should specify a callback function (*func* parameter) when calling the **unweldq** utility. If the caller also specifies a synchronization queue (*protect\_q* parameter) , STREAMS acquires the synchronization associated with that queue when calling *func*. If the callback function is not a protected STREAMS utility, such as the **qenable** utility, the caller should always specify a *protect\_q* parameter. The caller can also use this parameter to synchronize callbacks with protected STREAMS utilities.

**Note:** The **stream.h** header file must be the last included header file of each source file using the stream library.

### Parameters

<i>q1</i>	Specifies the queue whose <b>q_next</b> pointer must be nulled.
<i>q2</i>	Specifies the queue that will be unwelded to <i>q1</i> .
<i>q3</i>	Specifies the second queue whose <b>q_next</b> pointer must be nulled. If the <b>unweldq</b> utility is used to unweld only one pair of queues, this parameter should be set to <b>NULL</b> .
<i>q4</i>	Specifies the queue that will be unwelded to <i>q3</i> .
<i>func</i>	Specifies an optional callback function that will execute when the unwelding operation has completed.
<i>arg</i>	Specifies the parameter for <i>func</i> .

*protect\_q* Specifies an optional synchronization queue that protects *func*.

## Return Values

Upon successful completion, **0** (zero) is returned. Otherwise, an error code is returned.

## Error Codes

The **unweldq** utility fails if the following is true:

Value	Description
<b>EAGAIN</b>	The weld record could not be allocated. The caller may try again.
<b>EINVAL</b>	One or more parameters are not valid.
<b>ENXIO</b>	The weld mechanism is not installed.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts*.

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

Welding Mechanism in *AIX Version 6.1 Communications Programming Concepts*.

The **weldq** utility.

---

## wantio Utility

### Purpose

Register direct I/O entry points with the stream head.

### Syntax

```
#include <sys/stream.h>
int wantio(queue_t *q, struct wantio *w)
```

### Parameters

*q* Pointer to the **queue** structure.  
*w* Pointer to the **wantio** structure.

### Description

The **wantio** STREAMS routine can be used by a STREAMS module or driver to register input/output (read/write/select) entry points with the stream head. The stream head then calls these entry points directly, by-passing all normal STREAMS processing, when an I/O request is detected. This service may be useful to increase STREAMS performance in cases where normal module processing is not required or where STREAMS processing is to be performed outside of this operating system.

STREAMS modules and drivers should precede a **wantio** call by sending a high priority M\_LETSPLAY message upstream. The M\_LETSPLAY message format is a message block containing an integer followed by a pointer to the write queue of the module or driver originating the M\_LETSPLAY message. The integer counts the number of modules that can permit direct I/O. Each module passes this message to its neighbor after incrementing the count if direct I/O is possible. When this message reaches the stream head, the stream head compares the count field with the number of modules and drivers in the stream. If the count is not equal to the number of modules, then a M\_DONTPLAY message is sent downstream

indicating direct I/O will not be permitted on the stream. If the count is equal, then queued messages are cleared by sending them downstream as M\_BACKWASH messages. When all messages are cleared, then an M\_BACKDONE message is sent downstream. This process starts at the stream head and is repeated in every module in the stream. Modules will wait to receive an M\_BACKDONE message from upstream. Upon receipt of this message, the module will send all queued data downstream as M\_BACKWASH messages. When all data is cleared, the module will send an M\_BACKDONE message to its downstream neighbor indicating that all data has been cleared from the stream to this point. wantio registration is cleared from a stream by issuing a **wantio** call with a NULL pointer to the wantio structure.

Multiprocessor serialization is the responsibility of the driver or module requesting direct I/O. The stream head acquires no STREAMS locks before calling the wantio entry point.

Currently, the write entry point of the **wantio** structure is ignored.

## Return Values

Returns 0 always.

## Related Information

The **wantmsg** utility.

The **queue** and **wantio** structures in **/usr/include/sys/stream.h**.

---

## wantmsg Utility

### Purpose

Allows a STREAMS message to bypass a STREAMS module if the module is not interested in the message.

### Syntax

```
int wantmsg(q, f)
queue_t * q;
int (*f)();
```

### Description

The **wantmsg** utility allows a STREAMS message to bypass a STREAMS module if the module is not interested in the message, resulting in performance improvements.

The module registers filter functions with the read and write queues of the module with the **wantmsg** utility. A filter function takes as input a message pointer and returns 1 if the respective queue is interested in receiving the message. Otherwise it returns 0. The **putnext** and **qreply** subroutines call a queue's filter function before putting a message on that queue. If the filter function returns 1, then **putnext** or **qreply** put the message on that queue. Otherwise, **putnext** or **qreply** bypass the module by putting the message on the next module's queue.

The filter functions must be defined so that a message bypasses a module only when the module does not need to see the message.

The **wantmsg** utility cannot be used if the module has a service routine associated with the queue specified by the *q* parameter. If **wantmsg** is called for a module that has a service routine associated with *q*, **wantmsg** returns a value of 0 without registering the filter function with *q*.

This utility is part of STREAMS Kernel Extensions.

## Parameters

- q* Specifies the read or write queue to which the filter function is to be registered.  
*f* Specifies the module's filter function that is called at the **putnext** or **qreply** time.

## Return Values

Upon successful completion, the **wantmsg** utility returns a 1, indicating that the filter function specified by the *f* parameter has been registered for the queue specified by the *q* parameter. In this case, the filter function is called from **putnext** or **qreply**. The **wantmsg** utility returns a value of 0 if the module has a service routine associated with the queue *q*, indicating that the filter function is not registered with *q*.

## Example

```
wantmsg(q, tioc_is_r_interesting);
        wantmsg(WR(q), tioc_is_w_interesting);

/*
 * read queue filter function.
 * queue is only interested in IOCNAK, IOCACK, and
 * CTL messages.
 */

static int
tioc_is_r_interesting(mblk_t *mp)
{
    if (mp->b_datap->db_type == M_DATA)
        /* fast path for data messages */
        return 0;
    else if (mp->b_datap->db_type == M_IOCNAK ||
             mp->b_datap->db_type == M_IOCACK ||
             mp->b_datap->db_type == M_CTL)
        return 1;
    else
        return 0;
}

/*
 * write queue filter function.
 * queue is only interested in IOCTL and IOCADATA
 * messages.
 */

static int
tioc_is_w_interesting(mblk_t *mp)
{
    if (mp->b_datap->db_type == M_DATA)
        /* fast path for data messages */
        return 0;
    else if (mp->b_datap->db_type == M_IOCTL ||
             mp->b_datap->db_type == M_IOCADATA)
        return 1;
    else
        return 0;
}
```

## Related Information

The **putnext** utility, the **qreply** utility.

List of Streams Programming References, STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## weldq Utility

### Purpose

Establishes an uni-directional connection between STREAMS queues.

### Syntax

```
#include <sys/stream.h>
```

```
int weldq ( q1, q2, q3, q4, func, arg, protect_q)
queue_t *q1;
queue_t *q2;
queue_t *q3;
queue_t *q4;
weld_fcn_t func;
weld_arg_t arg;
queue_t *protect_q;
```

### Description

The **weldq** utility establishes an uni-directional connection (weld connection) between two STREAMS queues (*q1* and *q2*). The **weldq** utility can be used to weld two pairs of queues in one call (*q1* and *q2*, *q3* and *q4*).

The welding operation is performed by changing the first queue's **q\_next** pointer to point to the second queue. The **weldq** utility does not actually perform the operation. Instead, it creates a welding request which STREAMS performs asynchronously. STREAMS acquires the appropriate synchronization queues before performing the operation.

Callers that need to know when the welding operation has actually taken place should specify a callback function (*func* parameter) when calling the **weldq** utility. If the caller also specifies a synchronization queue (*protect\_q* parameter), STREAMS acquires the synchronization associated with that queue when calling *func*. If the callback function is not a protected STREAMS utility, such as the **qenable** utility, the caller should always specify a *protect\_q* parameter. The caller can also use this parameter to synchronize callbacks with protected STREAMS utilities.

This utility is part of STREAMS Kernel Extensions.

**Note:** The **stream.h** header file must be the last included header file of each source file using the stream library.

### Parameters

<i>q1</i>	Specifies the queue whose <b>q_next</b> pointer must be modified.
<i>q2</i>	Specifies the queue that will be welded to <i>q1</i> .
<i>q3</i>	Specifies the second queue whose <b>q_next</b> pointer must be modified. If the <b>weldq</b> utility is used to weld only one pair of queues, this parameter should be set to <b>NULL</b> .
<i>q4</i>	Specifies the queue that will be welded to <i>q3</i> .
<i>func</i>	Specifies an optional callback function that will execute when the welding operation has completed.
<i>arg</i>	Specifies the parameter for <i>func</i> .
<i>protect_q</i>	Specifies an optional synchronization queue that protects <i>func</i> .

### Return Values

Upon successful completion, **0** (zero) is returned. Otherwise, an error code is returned.

## Error Codes

The **weldq** utility fails if the following is true:

Value	Description
<b>EAGAIN</b>	The weld record could not be allocated. The caller may try again.
<b>EINVAL</b>	One or more parameters are not valid.
<b>ENXIO</b>	The weld mechanism is not installed.

## Related Information

List of Streams Programming References in *AIX Version 6.1 Communications Programming Concepts*.

STREAMS Overview in *AIX Version 6.1 Communications Programming Concepts*.

Welding Mechanism in *AIX Version 6.1 Communications Programming Concepts*.

The **unweldq** utility.

---

## WR Utility

### Purpose

Retrieves a pointer to the write queue.

### Syntax

```
#define WR( q) ((q)+1)
```

### Description

The **WR** utility accepts a read queue pointer, the *q* parameter, as an argument and returns a pointer to the write queue for the same module.

This utility is part of STREAMS Kernel Extensions.

### Parameters

*q* Specifies the read queue.

## Related Information

The **OTHERQ** utility, **RD** utility.

List of Streams Programming References and Understanding STREAMS Messages in *AIX Version 6.1 Communications Programming Concepts*.

---

## xtiso STREAMS Driver

### Purpose

Provides access to sockets-based protocols to STREAMS applications.

## Description

The **xtiso** driver (X/Open Transport Interface (XTI) over Sockets) is a STREAMS-based pseudo-driver that provides a Transport Layer Interface (TLI) to the socket-based protocols. The only supported use of the **xtiso** driver is by the TLI and XTI libraries.

The TLI and XTI specifications do not describe the name of the transport provider and how to address local and remote hosts, two important items required for use.

The **xtiso** driver supports most of the protocols available through the socket interface. Each protocol has a **/dev** entry, which must be used as the *name* parameter in the **t\_open** subroutine. The currently supported names (as configured by the **strload** subroutine) are:

Name	Socket Equivalent
/dev/xti/unixdg	AF_UNIX, SOCK_DGRAM
/dev/xti/unixst	AF_UNIX, SOCK_STREAM
/dev/xti/udp	AF_INET, SOCK_DGRAM
/dev/xti/tcp	AF_INET, SOCK_STREAM

Each of these protocols has a **sockaddr** structure that is used to specify addresses. These structures are also used by the TLI and XTI functions that require host addresses. The **netbuf** structure associated with the address for a particular function should refer to one of the **sockaddr** structure types. For instance, the TCP socket protocol uses a **sockaddr\_in** structure; so a corresponding **netbuf** structure would be:

```
struct netbuf addr;
struct sockaddr_in sin;
/* initialize sockaddr here */
sin.sin_family = AF_INET;
sin.sin_port = 0;
sin.sin_addr.s_addr = inet_addr("127.0.0.1");
addr.maxlen = sizeof(sin);
addr.len = sizeof(sin);
addr.buf = (char *)&sin;
```

The XTI Stream always consists of a Stream head and the transport interface module, **timod**. Depending on the transport provider specified by the application, **timod** accesses either the STREAMS-based protocol stack natively or a socket-based protocol through the pseudo-driver, **xtiso**.

The XTI library, **libxti.a** assumes a STREAMS-based transport provider. The routines of this library perform various operations for sending transport Provider Interface, TPI, messages down the XTI streams to the transport provider and receives them back.

The transport interface module, **timod**, is a STREAMS module that completes the translation of the TPI messages in the downstream and upstream directions.

The **xtiso** driver is a pseudo-driver that acts as the transport provider for socket-based communications. It interprets back and forth between the the TPI messages it receives from upstream and the socket interface.

AIX also provides the transport interface read/write module, **tirdwr**, which applications can push on to the XTI/TLI Stream for accessing the socket layer with standard UNIX read and write calls.

This driver is part of STREAMS Kernel Extensions.

## Files

**/dev/xti/\***                      Contains names of supported protocols.

## Related Information

The **strload** command.

The **t\_bind** subroutine for Transport Layer Interface, **t\_connect** subroutine for Transport Layer Interface, **t\_open** subroutine for Transport Layer Interface.

The **t\_bind** subroutine for X/Open Transport Layer Interface, **t\_connect** subroutine for X/Open Transport Layer Interface, **t\_open** subroutine for X/Open Transport Layer Interface.

Internet Transport-Level Protocols in *Networks and communication management*.

*UNIX System V Release 4 Programmer's Guide: Networking Interfaces*.

Understanding STREAMS Drivers and Modules in *AIX Version 6.1 Communications Programming Concepts*.

---

## t\_accept Subroutine for X/Open Transport Interface

### Purpose

Accept a connect request.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_accept (fd, resfd, call)
int fd;
int resfd;
const struct t_call *call;
```

### Description

The **t\_accept** subroutine is issued by a transport user to accept a command request. A transport user may accept a connection on either the same local transport endpoint or on an endpoint different than the one on which the connect indication arrived.

Before the connection can be accepted on the same endpoint, the user must have responded to any previous connect indications received on that transport endpoint via the **t\_accept** subroutine or the **t\_snddis** subroutine. Otherwise, the **t\_accept** subroutine will fail and set **t\_errno** to **TINDOUT**.

If a different transport endpoint is specified, the user may or may not choose to bind the endpoint before the **t\_accept** subroutine is issued. If the endpoint is not bound prior to the **t\_accept** subroutine, the transport provider will automatically bind the endpoint to the same protocol address specified in the *fd* parameter. If the transport user chooses to bind the endpoint, it must be bound to a protocol address with a *qlen* field of zero (see the **t\_bind** subroutine) and must be in the **T\_IDLE** state before the **t\_accept** subroutine is issued.

The call to the **t\_accept** subroutine fails with **t\_errno** set to **TLOOK** if there are indications (for example, connect or disconnect) waiting to be received on the endpoint specified by the *fd* parameter.

The value specified in the *udata* field enables the called transport user to send user data to the caller. The amount of user data sent must not exceed the limits supported by the transport provider. This limit is

specified in the *connect* field of the **t\_info** structure of the **t\_open** or **t\_getinfo** subroutines. If the *len* field of *udata* is zero, no data is sent to the caller. All the *maxlen* fields are meaningless.

When the user does not indicate any option, it is assumed that the connection is to be accepted unconditionally. The transport provider may choose options other than the defaults to ensure that the connection is accepted successfully.

There may be transport provider-specific restrictions on address binding. See Appendix A, ISO Transport Protocol Information and Appendix B, Internet Protocol-specific Information.

Some transport providers do not differentiate between a connect indication and the connection itself. If the connection has already been established after a successful return of the **t\_listen** subroutine, the **t\_accept** subroutine will assign the existing connection to the transport endpoint specified by *resfd* (see Appendix B, Internet Protocol-specific Information).

## Parameters

*fd* Identifies the local transport endpoint where the connect indication arrived.  
*resfd* Specifies the local transport endpoint where the connection is to be established.  
*call* Contains information required by the transport provider to complete the connection. The *call* parameter points to a **t\_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

The fields within the structure have the following meanings:

*addr* Specifies the protocol address of the calling transport user. The address of the caller may be null (length zero). When this field is not null, the field may be optionally checked by the X/Open Transport Interface.

*opt* Indicates any options associated with the connection.

*udata* Points to any user data to be returned to the caller.

*sequence* Specifies the value returned by the **t\_listen** subroutine which uniquely associates the response with a previously received connect indication.

## Valid States

**fd**: **T\_INCON**  
**resfd** (*Fd* != *resfd*): **T\_IDLE**

## Return Values

**0** Successful completion.  
**-1** Unsuccessful completion, **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to accept a connection on the responding transport endpoint or to use the specified options.
<b>TBADADDR</b>	The specified protocol address was in an incorrect format or contained illegal information.

Value	Description
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The file descriptor <i>fd</i> or <i>resfd</i> does not refer to a transport endpoint.
<b>TBADOPT</b>	The specified options were in an incorrect format or contained illegal information.
<b>TBADSEQ</b>	An invalid sequence number was specified.
<b>TINDOUT</b>	The subroutine was called with the same endpoint, but there are outstanding connection indications on the endpoint. Those other connection indications must be handled either by rejecting them via the <b>t_snddis</b> subroutine or accepting them on a different endpoint via the <b>t_accept</b> subroutine.
<b>TLOOK</b>	An asynchronous event has occurred on the transport endpoint referenced by <i>fd</i> and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was called in the wrong sequence on the transport endpoint referenced by <i>fd</i> , or the transport endpoint referred to by <i>resfd</i> is not in the appropriate state.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface( <b>t_errno</b> ).
<b>TPROVMISMATCH</b>	The file descriptors <i>fd</i> and <i>resfd</i> do not refer to the same transport provider.
<b>TRESADDR</b>	This transport provider requires both <i>fd</i> and <i>resfd</i> to be bound to the same address. This error results if they are not.
<b>TRESQLEN</b>	The endpoint referenced by <i>resfd</i> (where <i>resfd</i> is a different transport endpoint) was bound to a protocol address with a <i>qlen</i> field value that is greater than zero.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_connect** subroutine, **t\_getstate** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_rcvconnect** subroutine.

---

## t\_alloc Subroutine for X/Open Transport Interface

### Purpose

Allocate a library structure.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
```

```
void *t_alloc (
int fd
int struct_type,
int fields)
```

### Description

The **t\_alloc** subroutine dynamically allocates memory for the various transport function parameter structures. This subroutine allocates memory for the specified structure, and also allocates memory for buffers referenced by the structure.

Use of the **t\_alloc** subroutine to allocate structures helps ensure the compatibility of user programs with future releases of the transport interface functions.

## Parameters

*fd* Specifies the transport endpoint through which the newly allocated structure will be passed.  
*struc\_type* Specifies the structure to be allocated. The possible values are:

**T\_BIND**  
struct t\_bind  
**T\_CALL**  
struct t\_call  
**T\_OPTMGMT**  
struct t\_optmgmt  
**T\_DIS** struct t\_discon  
**T\_UNITDATA**  
struct t\_unitdata  
**T\_UDERROR**  
struct t\_uderr  
**T\_INFO**  
struct t\_info

Each of these structures may subsequently be used as a parameter to one or more transport functions. Each of the above structures, except **T\_INFO**, contains at least one field of the **struct netbuf** type. For each field of this type, the user may specify that the buffer for that field should be allocated as well. The length of the buffer allocated will be equal to or greater than the appropriate size as returned in the *info* parameter of the **t\_open** or **t\_getinfo** subroutines.  
*fields* Specifies whether the buffer should be allocated for each field type. The *fields* parameter specifies which buffers to allocate, where the parameter is the bitwise-OR of any of the following:

**T\_ADDR** The *addr* field of the **t\_bind**, **t\_call**, **t\_unitdata** or **t\_uderr** structures.  
**T\_OPT** The *opt* field of the **t\_optmgmt**, **t\_call**, **t\_unitdata** or **t\_uderr** structures.  
**T\_UDATA** The *udata* field of the **t\_call**, **t\_discon** or **t\_unitdata** structures.  
**T\_ALL** All relevant fields of the given structure. Fields which are not supported by the transport provider specified by the *fd* parameter are not allocated.

For each relevant field specified in the *fields* parameter, the **t\_alloc** subroutine allocates memory for the buffer associated with the field and initializes the *len* field to zero and initializes the *buf* pointer and *maxlen* field accordingly. Irrelevant or unknown values passed in fields are ignored. The length of the buffer allocated is based on the same size information returned to the user on a call to the **t\_open** and **t\_getinfo** subroutines. Thus, the *fd* parameter must refer to the transport endpoint through which the newly allocated structure is passed so that the appropriate size information is accessed. If the size value associated with any specified field is -1 or -2, (see the **t\_open** or **t\_getinfo** subroutines), the **t\_alloc** subroutine is unable to determine the size of the buffer to allocate and fails, setting **t\_errno** to **TSYSERR** and *errno* to **EINVAL**. For any field not specified in *fields*, *buf* will be set to the null pointer and *len* and *maxlen* will be set to zero.

## Valid States

ALL - apart from T\_UNINIT.

## Return Values

On successful completion, the **t\_alloc** subroutine returns a pointer to the newly allocated structure. On failure, a null pointer is returned.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TSYSERR</b>	A system error has occurred during execution of this function.
<b>TNOSTRUCTYPE</b>	Unsupported structure type ( <i>struct_type</i> ) requested. This can include a request for a structure type which is inconsistent with the transport provider type specified, for example, connection-oriented or connectionless.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_free** subroutine, **t\_getinfo** subroutine, **t\_open** subroutine.

---

## t\_bind Subroutine for X/Open Transport Interface

### Purpose

Bind an address to a transport endpoint.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
```

```
int t_bind (fd, req, ret)
    int fd;
    const struct t_bind *req;
    struct t_bind *ret;
```

### Description

The **t\_bind** subroutine associates a protocol address with the transport endpoint specified by the *fd* parameter and activates that transport endpoint. In connection mode, the transport provider may begin enqueueing incoming connect indications or servicing a connection request on the transport endpoint. In connectionless mode, the transport user may send or receive data units through the transport endpoint.

The *req* and *ret* parameters point to a **t\_bind** structure containing the following members:

```
struct netbuf addr;
unsigned qlen;
```

Within this structure, the fields have the following meaning:

Field	Description
<i>addr</i>	Specifies a protocol address.
<i>qlen</i>	Indicates the maximum number of outstanding connect indications.

If the requested address is not available, the **t\_bind** subroutine returns **-1** with **t\_errno** set as appropriate. If no address is specified in the *req* parameter, (that is, the *len* field of the *addr* field in the *req* parameter is zero or the *req* parameter is NULL), the transport provider assigns an appropriate address to be bound,

and returns that address in the *addr* field of the *ret* parameter. If the transport provider could not allocate an address, the **t\_bind** subroutine fails with **t\_errno** set to **TNOADDR**.

The *qlen* field has meaning only when initializing a connection-mode service. This field specifies the number of outstanding connect indications that the transport provider should support for the given transport endpoint. An outstanding connect indication is one that has been passed to the transport user by the transport provider but which has not been accepted or rejected. A *qlen* field value of greater than zero is only meaningful when issued by a passive transport user that expects other users to call it. The value of the *qlen* field is negotiated by the transport provider and may be changed if the transport provider cannot support the specified number of outstanding connect indications. However, this value of the *qlen* field is never negotiated from a requested value greater than zero to zero. This is a requirement on transport providers. See "Implementation Specifics" for more information. On return, the *qlen* field in the *ret* parameter contains the negotiated value.

The requirement that the value of the *qlen* field never be negotiated from a requested value greater than zero to zero implies that transport providers, rather than the X/Open Transport Interface implementation itself, accept this restriction.

A transport provider may not allow an explicit binding of more than one transport endpoint to the same protocol address, although it allows more than one connection to be accepted for the same protocol address. To ensure portability, it is, therefore, recommended not to bind transport endpoints that are used as responding endpoints, (those specified in the *resfd* parameter), in a call to the **t\_accept** subroutine, if the responding address is to be the same as the called address.

## Parameters

*fd* Specifies the transport endpoint. If the *fd* parameter refers to a connection-mode service, this function allows more than one transport endpoint to be bound to the same protocol address. However, the transport provider must also support this capability and it is not possible to bind more than one protocol address to the same transport endpoint. If a user binds more than one transport endpoint to the same protocol address, only one endpoint can be used to listen for connect indications associated with that protocol address. In other words, only one **t\_bind** for a given protocol address may specify a *qlen* field value greater than zero. In this way, the transport provider can identify which transport endpoint should be notified of an incoming connect indication. If a user attempts to bind a protocol address to a second transport endpoint with a *qlen* field value greater than zero, **t\_bind** will return -1 and set **t\_errno** to **TADDRBUSY**. When a user accepts a connection on the transport endpoint that is being used as the listening endpoint, the bound protocol address will be found to be busy for the duration of the connection, until a **t\_unbind** or **t\_close** call has been issued. No other transport endpoints may be bound for listening on that same protocol address while that initial listening endpoint is active (in the data transfer phase or in the **T\_IDLE** state). This will prevent more than one transport endpoint bound to the same protocol address from accepting connect indications.

If the *fd* parameter refers to a connectionless-mode service, only one endpoint may be associated with a protocol address. If a user attempts to bind a second transport endpoint to an already bound protocol address, **t\_bind** will return -1 and set **t\_errno** to **TADDRBUSY**.

*req* Specifies the address to be bound to the given transport endpoint. The *req* parameter is used to request that an address, represented by the **netbuf** structure, be bound to the given transport endpoint. The **netbuf** structure is described in the **xti.h** file. In the *req* parameter, the **netbuf** structure *addr* fields have the following meanings:

*buf* Points to the address buffer.

*len* Specifies the number of bytes in the address.

*maxlen* Has no meaning for the *req* parameter.

The *req* parameter may be a null pointer if the user does not specify an address to be bound. Here, the value of the *qlen* field is assumed to be zero, and the transport provider assigns an address to the transport endpoint. Similarly, the *ret* parameter may be a null pointer if the user does not care what address was bound by the provider and is not interested in the negotiated value of the *qlen* field. It is valid to set the *req* and *ret* parameters to the null pointer for the same call, in which case the provider chooses the address to bind to the transport endpoint and does not return that information to the user.

*ret* Specifies the maximum size of the address buffer. On return, the *ret* parameter contains the address that the transport provider actually bound to the transport endpoint; this is the same as the address specified by the user in the *req* parameter. In the *ret* parameter, the **netbuf** structure fields have the following meanings:

*buf* Points to the buffer where the address is to be placed. On return, this points to the bound address.

*len* Specifies the number of bytes in the bound address on return.

*maxlen* Specifies the the maximum size of the address buffer. If the value of the *maxlen* field is not large enough to hold the returned address, an error will result.

## Valid States

T\_UNBIND.

## Return Values

0 Successful completion.

-1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
TACCES	The user does not have permission to use the specified address.
TADDRBUSY	The requested address is in use.
TBADADDR	The specified protocol address was in an incorrect format or contained illegal information.
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBUFOVLW	The number of bytes allowed for an incoming argument ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that argument. The provider's state will change to <b>T_IDLE</b> and the information to be returned in <i>ret</i> will be discarded.
TNOADDR	The transport provider could not allocate an address.
TOUTSTATE	The function was issued in the wrong sequence.
TPROTO	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
TSYSERR	A system error has occurred during execution of this function.

## Related Information

The **t\_alloc** subroutine, **t\_close** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_unbind** subroutine.

---

## t\_close Subroutine for X/Open Transport Interface

### Purpose

Close a transport endpoint.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_close (fd)
int fd;
```

### Description

The **t\_close** subroutine informs the transport provider that the user is finished with the transport endpoint specified by the *fd* parameter and frees any local library resources associated with the endpoint. In addition, the **t\_close** subroutine closes the file associated with the transport endpoint.

The **t\_close** subroutine should be called from the **T\_UNBND** state (see the **t\_getstate** subroutine). However, this subroutine does not check state information, so it may be called from any state to close a transport endpoint. If this occurs, the local library resources associated with the endpoint will be freed automatically. In addition, the **close** subroutine is issued for that file descriptor. The **close** subroutine is abortive if there are no other descriptors in this process or if there are no other descriptors in another process which references the transport endpoint, and in this case, will break any transport connection that may be associated with that endpoint.

A **t\_close** subroutine issued on a connection endpoint may cause data previously sent, or data not yet received, to be lost. It is the responsibility of the transport user to ensure that data is received by the remote peer.

### Parameter

*fd*      Specifies the transport endpoint to be closed.

### Valid States

ALL - apart from T\_UNINIT.

### Return Values

0          Successful completion.  
-1        **t\_errno** is set to indicate an error.

### Errors

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The `t_getstate` subroutine, `t_open` subroutine, `t_unbind` subroutine.

---

## t\_connect Subroutine for X/Open Transport Interface

### Purpose

Establish a connection with another transport user.

### Library

X/Open Transport Interface Library (`libxti.a`)

### Syntax

```
#include <xti.h>

int t_connect (fd, sndcall, rcvcall)
    int fd;
    const struct t_call *sndcall;
    struct t_call *rcvcall;
```

### Description

The `t_connect` subroutine enables a transport user to request a connection to the specified destination transport user. This subroutine can only be issued in the **T\_IDLE** state.

The `sndcall` and `rcvcall` parameters both point to a `t_call` structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

In the `sndcall` parameter, the fields of the structure have the following meanings:

Field	Description
<code>addr</code>	Specifies the protocol address of the destination transport user.
<code>opt</code>	Presents any protocol-specific information that might be needed by the transport provider.
<code>sequence</code>	Has no meaning for this subroutine.
<code>udata</code>	Points to optional user data that may be passed to the destination transport user during connection establishment.

On return, the fields of the structure pointed to by the `rcvcall` parameter have the following meanings:

Field	Description
<code>addr</code>	Specifies the protocol address associated with the responding transport endpoint.
<code>opt</code>	Represents any protocol-specific information associated with the connection.
<code>sequence</code>	Has no meaning for this subroutine.
<code>udata</code>	Points to optional user data that may be returned by the destination transport user during connection establishment.

The `opt` field permits users to define the options that may be passed to the transport provider. These options are specific to the underlying protocol of the transport provider and are described for ISO and TCP protocols in Appendix A, ISO Transport Protocol Information, Appendix B, Internet Protocol-specific Information and Appendix F, Headers and Definitions. The user may choose not to negotiate protocol options by setting the `len` field of `opt` to zero. In this case, the provider may use default options.

If used, the value of the *opt.buf* field of the *sndcall* parameter **netbuf** structure must point to a buffer with the corresponding options; the *maxlen* and *buf* values of the *addr* and *opt* fields of the *rcvcall* parameter **netbuf** structure must be set before the call.

The *udata* field of the structure enables the caller to pass user data to the destination transport user and receive user data from the destination user during connection establishment. However, the amount of user data must not exceed the limits supported by the transport provider as returned in the *connect* field of the *info* parameter of the **t\_open** or **t\_getinfo** subroutines. If the value of *udata.len* field is zero in the *sndcall* parameter **netbuf** structure, no data will be sent to the destination transport user.

On return, the *addr*, *opt*, and *udata* fields of *rcvcall* are updated to reflect values associated with the connection. Thus, the *maxlen* value of each field must be set before issuing this subroutine to indicate the maximum size of the buffer for each. However, the value of the *rcvcall* parameter may be a null pointer, in which case no information is given to the user on return from the **t\_connect** subroutine.

By default, the **t\_connect** subroutine executes in synchronous mode, and waits for the destination user's response before returning control to the local user. A successful return (for example, return value of zero) indicates that the requested connection has been established. However, if **O\_NONBLOCK** is set via the **t\_open** subroutine or the *fcntl* parameter, the **t\_connect** subroutine executes in asynchronous mode. In this case, the call will not wait for the remote user's response, but returns control immediately to the local user and returns -1 with **t\_errno** set to **TNODATA** to indicate that the connection has not yet been established. In this way, the subroutine initiates the connection establishment procedure by sending a connect request to the destination transport user. The **t\_rcvconnect** subroutine is used in conjunction with the **t\_connect** subroutine to determine the status of the requested connection.

When a synchronous **t\_connect** call is interrupted by the arrival of a signal, the state of the corresponding transport endpoint is **T\_OUTCON**, allowing a further call to either the **t\_rcvconnect**, **t\_rcvdis** or **t\_snddis** subroutines.

## Parameters

<i>fd</i>	Identifies the local transport endpoint where communication will be established.
<i>sndcall</i>	Specifies information needed by the transport provider to establish a connection.
<i>rcvcall</i>	Specifies information associated with the newly established connection.

## Valid States

**T\_IDLE**.

## Return Values

0	Successful completion.
-1	<b>t_errno</b> is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to use the specified address or options.
<b>TADDRBUSY</b>	This transport provider does not support multiple connections with the same local and remote addresses. This error indicates that a connection already exists.
<b>TBADADDR</b>	The specified protocol address was in an incorrect format or contained illegal information.
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.

Value	Description
<b>TBADOPT</b>	The specified protocol options were in an incorrect format or contained illegal information.
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming parameter ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that argument. If executed in synchronous mode, the provider's state, as seen by the user, changes to <b>T_DAXAXFER</b> , and the information to be returned in the <i>rcvcall</i> parameter is discarded.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.,
<b>TNODATA</b>	<b>O_NONBLOCK</b> was set, so the subroutine successfully initiated the connection establishment procedure, but did not wait for a response from the remote user.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_accept** subroutine, **t\_alloc** subroutine, **t\_getinfo** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_rcvconnect** subroutine.

---

## t\_error Subroutine for X/Open Transport Interface

### Purpose

Produce error message.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_error (
    const char *errmsg)
```

### Description

The **t\_error** subroutine produces a language-dependent message on the standard error output which describes the last error encountered during a call to a transport subroutine.

If the *errmsg* parameter is not a null pointer and the character pointed to by the *errmsg* parameter is not the null character, the error message is written as follows: the string pointed to by the *errmsg* parameter followed by a colon and a space and a standard error message string for the current error defined in **t\_errno**. If **t\_errno** has a value different from **TSYSERR**, the standard error message string is followed by a newline character. If, however, **t\_errno** is equal to **TSYSERR**, the **t\_errno** string is followed by the standard error message string for the current error defined in the *errno* global variable followed by a newline.

The language for error message strings written by the **t\_error** subroutine is implementation-defined. If it is in English, the error message string describing the value in **t\_errno** is identical to the comments following the **t\_errno** codes defined in the **xti.h** header file. The contents of the error message strings describing the value in the *errno* global variable are the same as those returned by the **strerror** subroutine with an parameter of *errno*.

The error number, **t\_errno**, is only set when an error occurs and it is not cleared on successful calls.

## Parameter

*errmsg* Specifies a user-supplied error message that gives the context to the error.

## Valid States

ALL - apart from T\_UNINIT.

## Return Values

Upon completion, a value of 0 is returned.

## Errors Codes

No errors are defined for the **t\_error** subroutine.

## Examples

If a **t\_connect** subroutine fails on transport endpoint fd2 because a bad address was given, the following call might follow the failure:

```
t_error("t_connect failed on fd2");
```

The diagnostic message to be printed would look like:

```
t_connect failed on fd2: incorrect addr format
```

where *incorrect addr format* identifies the specific error that occurred, and *t\_connect failed on fd2* tells the user which function failed on which transport endpoint.

## Related Information

The **strerror** subroutine, **t\_connect** subroutine.

---

## t\_free Subroutine for X/Open Transport Interface

### Purpose

Free a library structure.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_free (
    void *ptr;
    int struct_type)
```

### Description

The **t\_free** subroutine frees memory previously allocated by the **t\_alloc** subroutine. This subroutine frees memory for the specified structure and buffers referenced by the structure.

The **t\_free** subroutine checks the *addr*, *opt*, and *udata* fields of the given structure, as appropriate, and frees the buffers pointed to by the *buf* field of the **netbuf** structure. If *buf* is a null pointer, the **t\_free**

subroutine does not attempt to free memory. After all buffers are free, the **t\_free** subroutine frees the memory associated with the structure pointed to by the *ptr* parameter.

Undefined results occur if the *ptr* parameter or any of the *buf* pointers points to a block of memory that was not previously allocated by the **t\_alloc** subroutine.

## Parameters

*ptr* Points to one of the seven structure types described for the **t\_alloc** subroutine.  
*struct\_type* Identifies the type of the structure specified by the *ptr* parameter. The type can be one of the following:

<b>T_BIND</b>	struct t_bind
<b>T_CALL</b>	struct t_call
<b>T_OPTMGMT</b>	struct t_optmgmt
<b>T_DIS</b>	struct t_discon
<b>T_UNITDATA</b>	struct t_unitdata
<b>T_UDERROR</b>	struct t_uderr
<b>T_INFO</b>	struct t_info

Each of these structures may subsequently be used as a parameter to one or more transport functions.

## Valid States

ALL - apart from T\_UNINIT.

## Return Values

0	Successful completion.
-1	<b>t_errno</b> is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TSYSERR</b>	A system error has occurred during execution of this function.
<b>TNOSTRUCTYPE</b>	Unsupported <i>struct_type</i> parameter value requested.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_alloc** subroutine.

---

## t\_getinfo Subroutine for X/Open Transport Interface

### Purpose

Get protocol-specific service information.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_getinfo (fd, info)
int fd;
struct t_info *info;
```

### Description

The **t\_getinfo** subroutine returns the current characteristics of the underlying transport protocol and/or transport connection associated with the file descriptor specified by the *fd* parameter. The pointer specified by the *info* parameter returns the same information returned by the **t\_open** subroutine, although not necessarily precisely the same values. This subroutine enables a transport user to access this information during any phase of communication.

### Parameters

<i>fd</i>	Specifies the file descriptor.
<i>info</i>	Points to a <b>t_info</b> structure which contains the following members:
long addr;	/* max size of the transport protocol       */ /* address                                       */
long options;	/* max number of bytes of protocol-specific */ /* options                                     */
long tsdu;	/* max size of a transport service data     */ /* unit (TSDU)                               */
long etsdu;	/* max size of an expedited transport       */ /* service data unit (ETSDU)               */
long connect;	/* max amount of data allowed on connection */ /* establishment functions                  */
long discon;	/* max amount of data allowed on t_snddis   */ /* and t_rcvdis functions                   */
long servtype;	/* service type supported by the transport */ /* provider                                  */
long flags;	/* other info about the transport provider */

The values of the fields have the following meanings:

Field	Description
addr	A value greater than zero indicates the maximum size of a transport protocol address and a value of -2 specifies that the transport provider does not provide user access to transport protocol addresses.
options	A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of -2 specifies that the transport provider does not support options set by users.
tsdu	A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of zero specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a datastream with no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of a TSDU; and a value of -2 specifies that the transfer of normal data is not supported by the transport provider.

Field	Description
etsdu	A value greater than zero specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of an ETSDU; and a value of -2 specifies that the transfer of expedited data is not supported by the transport provider. Note that the semantics of expedited data may be quite different for different transport providers (see Appendix A, ISO Transport Protocol Information and Appendix B, Internet Protocol-specific Information) .
connect	A value greater than zero specifies the maximum amount of data that may be associated with connection establishment functions and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment functions.
discon	A value greater than zero specifies the maximum amount of data that may be associated with the <b>t_snddis</b> and <b>t_rcvdis</b> subroutines and a value of -2 specifies that the transport provider does not allow data to be sent with the abortive release functions.
servtype	This field specifies the service type supported by the transport provider on return. The possible values are:  <b>T_COTS</b> The transport provider supports a connection-mode service but does not support the optional orderly release facility.  <b>T_COTS_ORD</b> The transport provider supports a connection-mode service with the optional orderly release facility.  <b>T_CLTS</b> The transport provider supports a connectionless-mode service. For this service type, the <b>t_open</b> subroutine will return -2 for etsdu, connect and discon.
flags	This is a bit field used to specify other information about the transport provider. If the <b>T_SENDZERO</b> bit is set in flags, this indicates that the underlying transport provider supports the sending of zero-length TSDUs. See Appendix A, ISO Transport Protocol Information for a discussion of the separate issue of zero-length fragments within a TSDU.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the **t\_alloc** subroutine may be used to allocate these buffers. An error results if a transport user exceeds the allowed data size on any subroutine. The value of each field may change as a result of protocol option negotiation during connection establishment (the **t\_optmgmt** call has no affect on the values returned by the **t\_getinfo** subroutine). These values will only change from the values presented to the **t\_open** subroutine after the endpoint enters the **T\_DATAXFER** state.

## Valid States

ALL - apart from T\_UNINIT.

## Return Values

- 0 Successful completion.
- 1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

Value	Description
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_alloc** subroutine, **t\_open** subroutine.

---

## t\_getprotaddr Subroutine for X/Open Transport Interface

### Purpose

Get the protocol addresses.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_getprotaddr (fd, boundaddr, peeraddr)
int fd;
struct t_bind *boundaddr;
struct t_bind *peeraddr;
```

### Description

The **t\_getproaddr** subroutine returns local and remote protocol addresses currently associated with the transport endpoint specified by the *fd* parameter.

### Parameters

<i>fd</i>	Specifies the transport endpoint.
<i>boundaddr</i>	Specifies the local address to which the transport endpoint is to be bound. The <i>boundaddr</i> parameter has the following fields: <ul style="list-style-type: none"> <li><i>maxlen</i> Specifies the maximum size of the address buffer.</li> <li><i>buf</i> Points to the buffer where the address is to be placed. On return, the <i>buf</i> field of <i>boundaddr</i> points to the address, if any, currently bound to <i>fd</i>.</li> <li><i>len</i> Specifies the length of the address. If the transport endpoint is in the <b>T_UNBND</b> state, zero is returned in the <i>len</i> field of <i>boundaddr</i>.</li> </ul>
<i>peeraddr</i>	Specifies the remote protocol address associated with the transport endpoint. <ul style="list-style-type: none"> <li><i>maxlen</i> Specifies the maximum size of the address buffer.</li> <li><i>buf</i> Points to the address, if any, currently connected to <i>fd</i>.</li> <li><i>len</i> Specifies the length of the address. If the transport endpoint is not in the <b>T_DATAXFER</b> state, zero is returned in the <i>len</i> field of <i>peeraddr</i>.</li> </ul>

### Valid States

ALL - apart from **T\_UNINIT**.

## Return Values

- 0 Successful completion.
- 1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBUFOVIEW</b>	The number of bytes allocated for an incoming parameter ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that parameter.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_bind** subroutine.

---

## t\_getstate Subroutine for X/Open Transport Interface

### Purpose

Get the current state.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_getstate (fd)
int fd;
```

### Description

The **t\_getstate** subroutine returns the current state of the provider associated with the transport endpoint specified by the *fd* parameter.

### Parameter

*fd* Specifies the transport endpoint.

### Valid States

ALL - apart from T\_UNINIT.

### Return Values

- 0 Successful completion.

-1 **t\_errno** is set to indicate an error. The current state is one of the following:

**T\_UNBND**

Unbound

**T\_IDLE**

Idle

**T\_OUTCON**

Outgoing connection pending

**T\_INCON**

Incoming connection pending

**T\_DATAXFER**

Data transfer

**T\_OUTREL**

Outgoing orderly release (waiting for an orderly release indication)

**T\_INREL**

Incoming orderly release (waiting to send an orderly release request)

If the provider is undergoing a state transition when the **t\_getstate** subroutine is called, the subroutine will fail.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TSTATECHNG</b>	The transport provider is undergoing a transient state change.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_open** subroutine.

---

## t\_listen Subroutine for X/Open Transport Interface

### Purpose

Listen for a connect indication.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_listen (fd, call)
int fd;
struct t_call *call;
```

### Description

The **t\_listen** subroutine listens for a connect request from a calling transport user.

By default, the **t\_listen** subroutine executes in synchronous mode and waits for a connect indication to arrive before returning to the user. However, if **O\_NONBLOCK** is set via the **t\_open** subroutine or with the **fcntl** subroutine (**F\_SETFL**), the **t\_listen** subroutine executes asynchronously, reducing to a poll for existing connect indications. If none are available, the subroutine returns -1 and sets **t\_errno** to **TNODATA**.

Some transport providers do not differentiate between a connect indication and the connection itself. If this is the case, a successful return of **t\_listen** indicates an existing connection (see Appendix B, Internet Protocol-specific Information).

## Parameters

*fd* Identifies the local transport endpoint where connect indications arrive.  
*call* Contains information describing the connect indication. The parameter *call* points to a **t\_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

In this structure, the fields have the following meanings:

*addr* Returns the protocol address of the calling transport user. This address is in a format usable in future calls to the **t\_connect** subroutine. Note, however that **t\_connect** may fail for other reasons, for example, **TADDRBUSY**.

*opt* Returns options associated with the connect request.

*udata* Returns any user data sent by the caller on the connect request.

*sequence*

A number that uniquely identifies the returned connect indication. The value of *sequence* enables the user to listen for multiple connect indications before responding to any of them.

Since this subroutine returns values for the *addr*, *opt* and *udata* fields of the *call* parameter, the *maxlen* field of each must be set before issuing the **t\_listen** subroutine to indicate the maximum size of the buffer for each.

## Valid States

**T\_IDLE**, **T\_INCON**.

## Return Values

0 Successful completion.  
-1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBADQLEN</b>	The <i>qlen</i> parameter of the endpoint referenced by the <i>fd</i> parameter is zero.
<b>TBODATA</b>	<b>O_NONBLOCK</b> was set, but no connect indications had been queued.

Value	Description
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming parameter ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that parameter. The provider's state, as seen by the user, changes to <b>T_INCON</b> , and the connect indication information to be returned in the <i>call</i> parameter is discarded. The value of the <i>sequence</i> parameter returned can be used to do a <b>t_snddis</b> .
<b>TLOOK</b>	An asynchronous event has occurred on the transport endpoint and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TQFULL</b>	The maximum number of outstanding indications has been reached for the endpoint referenced by the <i>fd</i> parameter.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **fcntl** subroutine, **t\_accept** subroutine, **t\_alloc** subroutine, **t\_bind** subroutine, **t\_connect** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine, **t\_rcvconnect** subroutine.

---

## t\_look Subroutine for X/Open Transport Interface

### Purpose

Look at the current event on a transport endpoint.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_look (fd)
int fd;
```

### Description

The **t\_look** subroutine returns the current event on the transport endpoint specified by the *fd* parameter. This subroutine enables a transport provider to notify a transport user of an asynchronous event when the user is calling subroutines in synchronous mode. Certain events require immediate notification of the user and are indicated by a specific error, **TLOOK**, on the current or next subroutine to be executed. Details on events which cause subroutines to fail, **T\_LOOK**, may be found in Section 4.6, Events and TLOOK Error Indication.

This subroutine also enables a transport user to poll a transport endpoint periodically for asynchronous events.

Additional functionality is provided through the Event Management (EM) interface.

### Parameter

*fd* Specifies the transport endpoint.

## Valid States

ALL - apart from **T\_UNINIT**.

## Return Values

Upon success, the **t\_look** subroutine returns a value that indicates which of the allowable events has occurred, or returns zero if no event exists. One of the following events is returned:

Event	Description
<b>T_LISTEN</b>	Connection indication received.
<b>T_CONNECT</b>	Connect confirmation received.
<b>T_DATA</b>	Normal data received.
<b>T_EXDATA</b>	Expedited data received.
<b>T_DISCONNECT</b>	Disconnect received.
<b>T_UDERR</b>	Datagram error indication.
<b>T_ORDREL</b>	Orderly release indication.
<b>T_GODATA</b>	Flow control restrictions on normal data flow that led to a <b>TFLOW</b> error have been lifted. Normal data may be sent again.
<b>T_GOEXDATA</b>	Flow control restrictions on expedited data flow that led to a <b>TFLOW</b> error have been lifted. Expedited data may be sent again.

On failure, -1 is returned and **t\_errno** is set to indicate the error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_open** subroutine, **t\_snd** subroutine, **t\_sndudata** subroutine.

---

## t\_open Subroutine for X/Open Transport Interface

### Purpose

Establish a transport endpoint.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

#include <fcntl.h>
int t_open (
    const char *name;
    int oflag;
    struct t_info *info)
```

## Description

The **t\_open** subroutine must be called as the first step in the initialization of a transport endpoint. This subroutine establishes a transport endpoint by supplying a transport provider identifier that indicates a particular transport provider (for example, transport protocol) and returning a file descriptor that identifies that endpoint.

This subroutine also returns various default characteristics of the underlying transport protocol by setting fields in the **t\_info** structure.

## Parameters

<i>name</i>	Points to a transport provider identifier.
<i>oflag</i>	Identifies any open flags (as in the <b>open</b> exec) . The <i>oflag</i> parameter is constructed from <b>O_RDWR</b> optionally bitwise inclusive-OR-ed with <b>O_NONBLOCK</b> . These flags are defined by the <b>fcntl.h</b> header file. The file descriptor returned by the <b>t_open</b> subroutine is used by all subsequent subroutines to identify the particular local transport endpoint.

*info* Points to a **t\_info** structure which contains the following members:

```
long addr;      /* max size of the transport protocol */
                /* address */
long options;   /* max number of bytes of
                /* protocol-specific options */
long tsdu;      /* max size of a transport service data
                /* unit (TSDU) */
long etsdu;     /* max size of an expedited transport
                /* service data unit (ETSDU) */
long connect;   /* max amount of data allowed on
                /* connection establishment subroutines */
long discon;    /* max amount of data allowed on
                /* t_snddis and t_rcvdis subroutines */
long servtype;  /* service type supported by the
                /* transport provider */
long flags;     /* other info about the transport provider */
```

The values of the fields have the following meanings:

**addr** A value greater than zero indicates the maximum size of a transport protocol address and a value of -2 specifies that the transport provider does not provide user access to transport protocol addresses.

**options** A value greater than zero indicates the maximum number of bytes of protocol-specific options supported by the provider, and a value of -2 specifies that the transport provider does not support user-settable options.

**tsdu** A value greater than zero specifies the maximum size of a transport service data unit (TSDU); a value of zero specifies that the transport provider does not support the concept of TSDU, although it does support the sending of a data stream with no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of a TSDU; and a value of -2 specifies that the transfer of normal data is not supported by the transport provider.

**etsdu** A value greater than zero specifies the maximum size of an expedited transport service data unit (ETSDU); a value of zero specifies that the transport provider does not support the concept of ETSDU, although it does support the sending of an expedited data stream with no logical boundaries preserved across a connection; a value of -1 specifies that there is no limit on the size of an ETSDU; and a value of -2 specifies that the transfer of expedited data is not supported by the transport provider. Note that the semantics of expedited data may be quite different for different transport providers.

**connect** A value greater than zero specifies the maximum amount of data that may be associated with connection establishment subroutines and a value of -2 specifies that the transport provider does not allow data to be sent with connection establishment subroutines.

**discon** A value greater than zero specifies the maximum amount of data that may be associated with the **t\_synddis** and **t\_rcvdis** subroutines and a value of -2 specifies that the transport provider does not allow data to be sent with the abortive release subroutines.

#### **servtype**

This field specifies the service type supported by the transport provider. The valid values on return are:

#### **T\_COTS**

The transport provider supports a connection-mode service but does not support the optional orderly release facility.

#### **T\_COTS\_ORD**

The transport provider supports a connection-mode service with the optional orderly release facility.

#### **T\_CLTS**

The transport provider supports a connectionless-mode service. For this service type, **t\_open** will return -2 for etsdu, connect and discon.

A single transport endpoint may support only one of the above services at one time.

**flags** This is a bit field used to specify other information about the transport provider. If the **T\_SENDCZERO** bit is set in flags, this indicates the underlying transport provider supports the sending of zero-length TSDUs.

If a transport user is concerned with protocol independence, the above sizes may be accessed to determine how large the buffers must be to hold each piece of information. Alternatively, the **t\_alloc** subroutine may be used to allocate these buffers. An error will result if a transport user exceeds the allowed data size on any subroutine.

If the *info* parameter is set to a null pointer by the transport user, no protocol information is returned by the **t\_open** subroutine.

## **Valid States**

**T\_UNINIT**

## **Return Values**

Valid file descriptor

-1

Successful completion.

**t\_errno** is set to indicate an error.

## **Error Codes**

On failure, **t\_errno** is set to one of the following:

<b>Value</b>	<b>Description</b>
<b>TBADFLAG</b>	An invalid flag is specified.
<b>TBADNAME</b>	Invalid transport provider name.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## **Related Information**

The **t\_open** subroutine.

---

## **t\_optmgmt Subroutine for X/Open Transport Interface**

### **Purpose**

Manage options for a transport endpoint.

### **Library**

X/Open Transport Interface Library (**libxti.a**)

### **Syntax**

```
#include <xti.h>
```

```
int t_optmgmt(  
    int fd,  
    const struct t_optmgmt *req,  
    struct t_optmgmt *ret)
```

### **Description**

The **t\_optmgmt** subroutine enables a transport user to retrieve, verify, or negotiate protocol options with the transport provider.

The *req* and *ret* parameters both point to a **t\_optmgmt** structure containing the following members:

```
struct netbuf opt;  
long flags;
```

Within this structure, the fields have the following meaning:

Field	Description
<i>opt</i>	Identifies protocol options. The options are represented by a <b>netbuf</b> structure in a manner similar to the address in the <b>t_bind</b> subroutine:
<i>len</i>	Specifies the number of bytes in the options and on return, specifies the number of bytes of options returned.
<i>buf</i>	Points to the options buffer. For the <i>ret</i> parameter, <i>buf</i> points to the buffer where the options are to be placed. Each option in the options buffer is of the form <b>struct t_opthdr</b> possibly followed by an option value. The fields of this structure and the values are:
<i>level</i>	Identifies the X/Open Transport Interface level or a protocol of the transport provider.
<i>name</i>	Identifies the option within the level.
<i>len</i>	Contains its total length, for example, the length of the option header <b>t_opthdr</b> plus the length of the option value. If <b>t_optmgmt</b> is called with the action <b>T_NEGOTIATE</b> set.
<i>status</i>	Contains information about the success or failure of a negotiation.
Each option in the input or output option buffer must start at a long-word boundary. The macro <b>OPT_NEXTHDR</b> ( <i>pbuf</i> , <i>buflen</i> , <i>poption</i> ) can be used for that purpose. The macro parameters are as follows:	
<i>pbuf</i>	Specifies a pointer to an option buffer <i>opt.buf</i> .
<i>buflen</i>	The length of the option buffer pointed to by <i>pbuf</i> .
<i>poption</i>	Points to the current option in the option buffer. <b>OPT_NEXTHDR</b> returns a pointer to the position of the next option or returns a null pointer if the option buffer is exhausted. The macro is helpful for writing and reading. See the <b>xti.h</b> header file for the exact definition of this structure.
If the transport user specifies several options on input, all options must address the same level.	
If any option in the options buffer does not indicate the same level as the first option, or the level specified is unsupported, then the <b>t_optmgmt</b> request fails with <b>TBADOPT</b> . If the error is detected, some options may have successfully negotiated. The transport user can check the current status by calling the <b>t_optmgmt</b> subroutine with the <b>T_CURRENT</b> flag set.	
<b>Note:</b> "The Use of Options" contains a detailed description about the use of options and should be read before using this subroutine.	
<i>maxlen</i>	Has no meaning for the <i>req</i> parameter, but must be set in the <i>ret</i> parameter to specify the maximum size of the options buffer. On return, <i>len</i> specifies the number of bytes of options returned. The value in <i>maxlen</i> has no meaning for the <i>req</i> argument,

Field	Description
<i>flags</i>	<p>Specifies the action to take with those options. The <i>flags</i> field of <i>req</i> must specify one of the following actions:</p> <p><b>T_CHECK</b></p> <p>This action enables the user to verify whether the options specified in the <i>req</i> parameter are supported by the transport provider. If an option is specified with no option value, (that is, it consists only of a <b>t_opthdr</b> structure), the option is returned with its <i>status</i> field set to one of the following:</p> <ul style="list-style-type: none"> <li>• <b>T_SUCCESS</b> - if it is supported.</li> <li>• <b>T_NOTSUPPORT</b> - if it is not or needs additional user privileges.</li> <li>• <b>T_READONLY</b> - if it is read-only (in the current X/Open Transport Interface state).</li> </ul> <p>No option value is returned. If an option is specified with an option value, the <i>status</i> field of the returned option has the same value, as if the user had tried to negotiate this value with <b>T_NEGOTIATE</b>. If the status is <b>T_SUCCESS</b>, <b>T_FAILURE</b>, <b>T_NOTSUPPORT</b>, or <b>T_READONLY</b>, the returned option value is the same as the one requested on input.</p> <p>The overall result of the option checks is returned in the <i>flags</i> field of the <b>netbuf</b> structure pointed to by the <i>ret</i> parameter. This field contains the worst single result of the option checks, where the rating is the same as for <b>T_NEGOTIATE</b>.</p> <p>Note, that no negotiation takes place. All currently effective option values remain unchanged.</p> <p><b>T_CURRENT</b></p> <p>This action enables the transport user to retrieve the currently effective option values. The user specifies the options of interest in the <i>opt</i> fields in the <b>netbuf</b> structure pointed to by the <i>req</i> parameter. The option values are irrelevant and will be ignored; it is sufficient to specify the <b>t_opthdr</b> part of an option only. The currently effective values are then returned in <i>opt</i> fields in the <b>netbuf</b> structure pointed to by the <i>ret</i> parameter.</p> <p>The <i>status</i> field returned is one of the following:</p> <ul style="list-style-type: none"> <li>• <b>T_NOTSUPPORT</b> if the protocol level does not support this option or the transport user illegally requested a privileged option.</li> <li>• <b>T_READONLY</b> if the option is read-only.</li> <li>• <b>T_SUCCESS</b> in all other cases.</li> </ul> <p>The overall result of the option checks is returned in the <i>flags</i> field of the <b>netbuf</b> structure pointed to by the <i>ret</i> parameter. This field contains the worst single result of the option checks, where the rating is the same as for <b>T_NEGOTIATE</b>.</p> <p>For each level, the <b>T_ALLOPT</b> option (see below) can be requested on input. All supported options of this level with their default values are then returned.</p>

Field	Description
<b>T_DEFAULT</b>	<p>This action enables the transport user to retrieve the default option values. The user specifies the options of interest in the <i>opt</i> fields in the <b>netbuf</b> structure pointed to by the <i>req</i> parameter. The option values are irrelevant and will be ignored; it is sufficient to specify the <b>t_opthdr</b> part of an option only. The default values are then returned in the <i>opt</i> field of the <b>netbuf</b> structure pointed to by the <i>ret</i> parameter.</p> <p>The <i>status</i> field returned is one of the following:</p> <ul style="list-style-type: none"> <li>• <b>T_NOTSUPPORT</b> if the protocol level does not support this option or the transport user illegally requested a privileged option.</li> <li>• <b>T_READONLY</b> if the option is read-only.</li> <li>• <b>T_SUCCESS</b> in all other cases.</li> </ul> <p>The overall result of the option checks is returned in the <i>flags</i> field of the <i>ret</i> parameter <b>netbuf</b> structure. This field contains the worst single result of the option checks, where the rating is the same as for <b>T_NEGOTIATE</b>.</p> <p>For each level, the <b>T_ALLOPT</b> option (see below) can be requested on input. All supported options of this level with their default values are then returned. In this case, the <i>maxlen</i> value of the <i>opt</i> field in the <i>ret</i> parameter <b>netbuf</b> structure must be given at least the value of the <i>options</i> field of the <i>info</i> parameter (see the <b>t_getinfo</b> or <b>t_open subroutines</b>) before the call.</p>
<b>T_NEGOTIATE</b>	<p>This action enables the transport user to negotiate option values. The user specifies the options of interest and their values in the buffer specified in the <i>req</i> parameter <b>netbuf</b> structure. The negotiated option values are returned in the buffer pointed to by the <i>opt</i> field of the <i>ret</i> parameter <b>netbuf</b> structure. The <i>status</i> field of each returned option is set to indicate the result of the negotiation. The value is one of the following:</p> <ul style="list-style-type: none"> <li>• <b>T_SUCCESS</b> if the proposed value was negotiated.</li> <li>• <b>T_PARTSUCCESS</b> if a degraded value was negotiated.</li> <li>• <b>T_FAILURE</b> is the negotiation failed (according to the negotiation rules).</li> <li>• <b>T_NOTSUPPORT</b> if the transport provider does not support this option or illegally requests negotiation of a privileged option</li> <li>• <b>T_READONLY</b> if modification of a read-only option was requested.</li> </ul> <p>If the status is <b>T_SUCCESS</b>, <b>T_FAILURE</b>, <b>T_NOTSUPPORT</b> or <b>T_READONLY</b>, the returned option value is the same as the one requested on input.</p> <p>The overall result of the negotiation is returned in the <i>flags</i> field of the <i>ret</i> parameter <b>netbuf</b> structure. This field contains the worst single result, whereby the rating is done according to the following order, where <b>T_NOTSUPPORT</b> is the worst result and <b>T_SUCCESS</b> is the best:</p> <ul style="list-style-type: none"> <li>• <b>T_NOTSUPPORT</b></li> <li>• <b>T_READONLY</b></li> <li>• <b>T_FAILURE</b></li> <li>• <b>T_PARTSUCCESS</b></li> <li>• <b>T_SUCCESS.</b></li> </ul> <p>For each level, the <b>T_ALLOPT</b> option (see below) can be requested on input. This option has no value and consists of a <b>t_opthdr</b> only. This input requests negotiation of all supported options of this level to their default values. The result is returned option by option in the <i>opt</i> field of the structure pointed to in the <i>ret</i> parameter. Depending on the state of the transport endpoint, not all requests to negotiate the default value may be successful.</p>

Field	Description
	<p>The <b>T_ALLOPT</b> option can only be used with the <b>t_optmgmt</b> structure and the actions <b>T_NEGOTIATE</b>, <b>T_DEFAULT</b> and <b>T_CURRENT</b>. This option can be used with any supported level and addresses all supported options of this level. The option has no value and consists of a <b>t_opthdr</b> only. Since only options of one level may be addressed in a <b>t_optmgmt</b> call, this option should not be requested together with other options. The subroutine returns as soon as this option has been processed.</p> <p>Options are independently processed in the order they appear in the input option buffer. If an option is multiply input, it depends on the implementation whether it is multiply output or whether it is returned only once.</p> <p>Transport providers may not be able to provide an interface capable of supporting <b>T_NEGOTIATE</b> and/or <b>T_CHECK</b> functionalities. When this is the case, the error <b>TNOTSUPPORT</b> is returned.</p> <p>The subroutine <b>t_optmgmt</b> may block under various circumstances and depending on the implementation. For example, the subroutine will block if the protocol addressed by the call resides on a separate controller. It may also block due to flow control constraints, if data previously sent across this transport endpoint has not yet been fully processed. If the subroutine is interrupted by a signal, the option negotiations that have been done so far may remain valid. The behavior of the subroutine is not changed if <b>O_NONBLOCK</b> is set.</p>

## Parameters

<i>fd</i>	Identifies a transport endpoint.
<i>req</i>	Requests a specific action of the provider.
<i>ret</i>	Returns options and flag values to the user.

## X/Open Transport Interface-Level Options

X/Open Transport Interface (XTI) level options are not specific for a particular transport provider. An XTI implementation supports none, all, or any subset of the options defined below. An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode, or if the bound transport endpoint identified by the *fd* parameter relates to specific transport providers.

The subsequent options are not association-related (see Chapter 5, The Use of Options) . They may be negotiated in all XTI states except **T\_UNINIT**.

The protocol level is **XTI\_GENERIC**. For this level, the following options are defined (the type of each option value is of type **unsigned long** unless otherwise indicated):

### XTI-Level Options

Option Name	Legal Option Value	Meaning
XTI_DEBUG (array of unsigned longs)	see text	enable debugging
XTI_LINGER (struct linger)	see text	linger on close if data is present
XTI_RCVBUF	size in octets	receive buffer size
XTI_RCVLOWAT	size in octets	receive low-water mark
XTI_SNDBUF0	size in octets	send buffer size
XTI_SNDLOWAT	size in octets	send low-water mark

A request for **XTI\_DEBUG** is an absolute requirement. A request to activate **XTI\_LINGER** is an absolute requirement; the timeout value to this option is not. **XTI\_RCVBUF**, **XTI\_RCVLOWAT**, **XTI\_SNDBUF** and **XTI\_SNDLOWAT** are not absolute requirements.

Option	Description
<b>XTI_DEBUG</b>	<p>Enables debugging. The values of this option are implementation-defined. Debugging is disabled if the option is specified with no value (for example, with an option header only).</p> <p>The system supplies utilities to process the traces. An implementation may also provide other means for debugging.</p>
<b>XTI_LINGER</b>	<p>Lingers the execution of a <b>t_close</b> subroutine or the <b>close</b> exec if send data is still queued in the send buffer. The option value specifies the linger period. If a <b>close</b> exec or <b>t_close</b> subroutine is issued and the send buffer is not empty, the system attempts to send the pending data within the linger period before closing the endpoint. Data still pending after the linger period has elapsed is discarded.</p> <p>Depending on the implementation, the <b>t_close</b> subroutine or <b>close</b> exec either, at a maximum, block the linger period, or immediately return, whereupon, at most, the system holds the connection in existence for the linger period.</p> <p>The option value consists of a <b>structure t_linger</b> declared as:</p> <pre> <b>struct t_linger</b> {     <b>long</b> <i>l_onoff</i>;     <b>long</b> <i>l_linger</i>; } </pre> <p>The fields of the structure and the legal values are:</p> <p><i>l_onoff</i> Switches the option on or off. The value <i>l_onoff</i> is an absolute requirement. The possible values are:</p> <p><b>T_NO</b> switch option off</p> <p><b>T_YES</b> activate option</p> <p><i>l_linger</i> Determines the linger period in seconds. The transport user can request the default value by setting the field to <b>T_UNSPEC</b>. The default timeout value depends on the underlying transport provider (it is often <b>T_INFINITE</b>). Legal values for this field are <b>T_UNSPEC</b>, <b>T_INFINITE</b> and all non-negative numbers.</p> <p>The <i>l_linger</i> value is not an absolute requirement. The implementation may place upper and lower limits to this value. Requests that fall short of the lower limit are negotiated to the lower limit.</p>
<b>XTI_RCVBUF</b>	<p><b>Note:</b> Note that this option does not linger the execution of the <b>t_snddis</b> subroutine.</p> <p>Adjusts the internal buffer size allocated for the receive buffer. The buffer size may be increased for high-volume connections, or decreased to limit the possible backlog of incoming data.</p> <p>This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.</p>
<b>XTI_RCVLOWAT</b>	<p>Legal values are all positive numbers.</p> <p>Sets a low-water mark in the receive buffer. The option value gives the minimal number of bytes that must have accumulated in the receive buffer before they become visible to the transport user. If and when the amount of accumulated receive data exceeds the low-water mark, a <b>T_DATA</b> event is created, an event mechanism (for example, the <b>poll</b> or <b>select</b> subroutines) indicates the data, and the data can be read by the <b>t_rcv</b> or <b>t_rcvudata</b> subroutines.</p> <p>This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.</p> <p>Legal values are all positive numbers.</p>

Option	Description
<b>XTI_SNDBUF</b>	Adjusts the internal buffer size allocated for the send buffer.  This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.
<b>XTI_SNDLOWAT</b>	Legal values are all positive numbers. Sets a low-water mark in the send buffer. The option value gives the minimal number of bytes that must have accumulated in the send buffer before they are sent.  This request is not an absolute requirement. The implementation may place upper and lower limits on the option value. Requests that fall short of the lower limit are negotiated to the lower limit.  Legal values are all positive numbers.

## Valid States

ALL - except from **T\_UNINIT**.

## Return Values

- 0** Successful completion.
- 1** **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TACCES</b>	The user does not have permission to negotiate the specified options.
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBADFLAG</b>	An invalid flag was specified.
<b>TBADOPT</b>	The specified options were in an incorrect format or contained illegal information.
<b>TBUFOVFLW</b>	The number of bytes allowed for an incoming argument ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that argument. The information to be returned in <i>ret</i> will be discarded.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_accept** subroutine, **t\_alloc** subroutine, **t\_connect** subroutine, **t\_getinfo** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_rcvconnect** subroutine.

---

## t\_rcv Subroutine for X/Open Transport Interface

### Purpose

Receive data or expedited data sent over a connection.

### Library

X/Open Transport Interface Library (**libxti.a**)

## Syntax

```
#include <xti.h>
```

```
int t_rcv (  
    int fd,  
    void *buf,  
    unsigned int nbytes,  
    int *flags)
```

## Description

The **t\_rcv** subroutine receives either normal or expedited data. By default, the **t\_rcv** subroutine operates in synchronous mode and waits for data to arrive if none is currently available. However, if **O\_NONBLOCK** is set via the **t\_open** subroutine or the *fcntl* parameter, the **t\_rcv** subroutine executes in asynchronous mode and fails if no data is available. (See the **TNODATA** error in "Error Codes" below.)

## Parameters

<i>fd</i>	Identifies the local transport endpoint through which data will arrive.
<i>buf</i>	Points to a receive buffer where user data will be placed.
<i>nbytes</i>	Specifies the size of the receive buffer.
<i>flags</i>	Specifies optional flags. This parameter may be set on return from the <b>t_rcv</b> subroutine. The possible values are:

### **T\_MORE**

If set, on return from the call, indicates that there is more data, and the current transport service data unit (TSDU) or expedited transport service data unit (ETSDU) must be received in multiple **t\_rcv** calls. In the asynchronous mode, the **T\_MORE** flag may be set on return from the **t\_rcv** call even when the number of bytes received is less than the size of the receive buffer specified. Each **t\_rcv** call with the **T\_MORE** flag set, indicates that another **t\_rcv** call must follow to get more data for the current TSDU. The end of the TSDU is identified by the return of a **t\_rcv** call with the **T\_MORE** flag not set. If the transport provider does not support the concept of a TSDU as indicated in the *info* parameter on return from the **t\_open** or **t\_getinfo** subroutines, the **T\_MORE** flag is not meaningful and should be ignored. If the *nbytes* parameter is greater than zero on the call to **t\_rcv**, **t\_rcv** returns 0 only if the end of a TSDU is being returned to the user.

### **T\_EXPEDITED**

If set, the data returned is expedited data. If the number of bytes of expedited data exceeds the value of the *nbytes* parameter, **t\_rcv** will set **T\_EXPEDITED** and **T\_MORE** on return from the initial call. Subsequent calls to retrieve the remaining ETSDU will have **T\_EXPEDITED** set on return. The end of the ETSDU is identified by the return of a **t\_rcv** call with the **T\_MORE** flag not set.

In synchronous mode, the only way to notify the user of the arrival of normal or expedited data is to issue this subroutine or check for the **T\_DATA** or **T\_EXDATA** events using the **t\_look** subroutine. Additionally, the process can arrange to be notified via the Event Management interface.

## Valid States

T\_DATAXFER, T\_OUTREL.

## Return Values

On successful completion, the **t\_rcv** subroutine returns the number of bytes received. Otherwise, it returns -1 on failure and **t\_errno** is set to indicate the error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNODATA</b>	<b>O_NONBLOCK</b> was set, but no data is currently available from the transport provider.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **fcntl** subroutine, **t\_getinfo** subroutine, **t\_look** subroutine, **t\_open** subroutine, **t\_snd** subroutine.

---

## t\_rcvconnect Subroutine for X/Open Transport Interface

### Purpose

Receive the confirmation from a connect request.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_rcvconnect (fd, call)
int fd;
struct t_call *call;
```

### Description

The **t\_rcvconnect** subroutine enables a calling transport user to determine the status of a previously sent connect request and is used in conjunction with the **t\_connect** subroutine to establish a connection in asynchronous mode. The connection is established on successful completion of this subroutine.

### Parameters

*fd*       Identifies the local transport endpoint where communication will be established.

*call* Contains information associated with the newly established connection. The *call* parameter points to a **t\_call** structure which contains the following members:

```
struct netbuf addr;  
struct netbuf opt;  
struct netbuf udata;  
int sequence;
```

The fields of the **t\_call** structure are:

*addr* Returns the protocol address associated with the responding transport endpoint.

*opt* Presents any options associated with the connection.

*udata* Points to optional user data that may be returned by the destination transport user during connection establishment.

*sequence*  
Has no meaning for this subroutine.

The *maxlen* field of each **t\_call** member must be set before issuing this subroutine to indicate the maximum size of the buffer for each. However, the value of the *call* parameter may be a null pointer, in which case no information is given to the user on return from the **t\_rcvconnect** subroutine. By default, the **t\_rcvconnect** subroutine executes in synchronous mode and waits for the connection to be established before returning. On return, the *addr*, *opt* and *udata* fields reflect values associated with the connection.

If **O\_NONBLOCK** is set (via the **t\_open** subroutine or **fcntl**), the **t\_rcvconnect** subroutine executes in asynchronous mode, and reduces to a poll for existing connect confirmations. If none are available, the **t\_rcvconnect** subroutine fails and returns immediately without waiting for the connection to be established. (See **TNODATA** in "Error Codes" below.) In this case, the **t\_rcvconnect** subroutine must be called again to complete the connection establishment phase and retrieve the information returned in the *call* parameter.

## Valid States

**T\_OUTCON**

## Return Values

0 Successful completion.  
-1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBUFOVFLW</b>	The number of bytes allocated for an incoming argument ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the value of that argument, and the connect information to be returned in <i>call</i> will be discarded. The provider's state, as seen by the user, will be changed to <b>T_DATAFER</b> .
<b>TLOOK</b>	An asynchronous event has occurred on the transport connection and requires immediate attention.
<b>TNODATA</b>	<b>O_NONBLOCK</b> was set, but a connect confirmation has not yet arrived.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_accept** subroutine, **t\_alloc** subroutine, **t\_bind** subroutine, **t\_connect** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_optmgmt** subroutine.

---

## t\_rcvdis Subroutine for X/Open Transport Interface

### Purpose

Retrieve information from disconnect.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_rcvdis (fd, discon)
int fd;
struct t_discon *discon;
```

### Description

The **t\_rcvdis** subroutine identifies the cause of a disconnect and retrieves any user data sent with the disconnect.

### Parameters

*fd* Identifies the local transport endpoint where the connection existed.  
*discon* Points to a **t\_discon** structure containing the following members:

```
struct netbuf udata;
int reason;
int sequence;
```

The **t\_discon** structure fields are:

*reason* Specifies the reason for the disconnect through a protocol-dependent reason code.

*udata* Identifies any user data that was sent with the disconnect.

*sequence*

May identify an outstanding connect indication with which the disconnect is associated. The *sequence* field is only meaningful when the **t\_rcvdis** subroutine is issued by a passive transport user who has executed one or more **t\_listen** subroutines and is processing the resulting connect indications. If a disconnect indication occurs, the *sequence* field can be used to identify which of the outstanding connect indications is associated with the disconnect.

If a user does not care if there is incoming data and does not need to know the value of the *reason* or *sequence* fields, the *discon* field value may be a null pointer and any user data associated with the disconnect will be discarded. However, if a user has retrieved more than one outstanding connect indication (via the **t\_listen** subroutine) and the *discon* field value is a null pointer, the user will be unable to identify with which connect indication the disconnect is associated.

### Valid States

**T\_DATAXFER**, **T\_OUTCON**, **T\_OUTREL**, **T\_INREL**, **T\_INCON**(*ocnt* > 0).

## Return Values

- 0 Successful completion.
- 1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBUOVFLW</b>	The number of bytes allocated for incoming data ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the data. If the <i>fd</i> parameter is a passive endpoint with <i>ocnt</i> > 1, it remains in state <b>T_INCON</b> ; otherwise, the endpoint state is set to <b>T_IDLE</b> .
<b>TNODIS</b>	No disconnect indication currently exists on the specified transport endpoint.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_alloc** subroutine, **t\_connect** subroutine, **t\_listen** subroutine, **t\_open** subroutine, **t\_snddis** subroutine.

---

## t\_rcvrel Subroutine for X/Open Transport Interface

### Purpose

Acknowledging receipt of an orderly release indication.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_rcvrel (fd)
int fd;
```

### Description

The **t\_rcvrel** subroutine is used to acknowledge receipt of an orderly release indication. After receipt of this indication, the user may not attempt to receive more data because such an attempt will block forever. However, the user may continue to send data over the connection if the **t\_sndrel** subroutine has not been called by the user. This function is an optional service of the transport provider, and is only supported if the transport provider returned the **T\_COTS\_ORD** service type on **t\_open** or **t\_getinfo** calls.

### Parameter

*fd* Identifies the local transport endpoint where the connection exists.

## Valid States

T\_DATAXFER, T\_OUTREL.

## Return Values

0        Successful completion.  
-1       **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
TBADF	The specified file descriptor does not refer to a transport endpoint.
TLOOK	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
TNOREL	No orderly release indication currently exists on the specified transport endpoint.
TNOTSUPPORT	This subroutine is not supported by the underlying transport provider.
TOUTSTATE	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
TPROTO	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
TSYSERR	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine, **t\_sndrel** subroutine.

---

## t\_rcvudata Subroutine for X/Open Transport Interface

### Purpose

Receive a data unit.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_rcvudata (fd, unitdata, flags)
int fd;
struct t_unitdata *unitdata;
int *flags;
```

### Description

The **t\_rcvudata** subroutine is used in connectionless mode to receive a data unit from another transport user.

By default, the **t\_rcvudata** subroutine operates in synchronous mode and waits for a data unit to arrive if none is currently available. However, if **O\_NONBLOCK** is set (via the **t\_open** subroutine or *fcntl*), the **t\_rcvudata** subroutine executes in asynchronous mode and fails if no data units are available.

If the buffer defined in the *udata* field of the *unitdata* parameter is not large enough to hold the current data unit, the buffer is filled and **T\_MORE** is set in the *flags* parameter on return to indicate that another **t\_rcvudata** subroutine should be called to retrieve the rest of the data unit. Subsequent calls to the **t\_rcvudata** subroutine return zero for the length and options until the full data unit is received.

## Parameters

<i>fd</i>	Identifies the local transport endpoint through which data will be received.
<i>unitdata</i>	Holds information associated with the received data unit. The <i>unitdata</i> parameter points to a <b>t_unitdata</b> structure containing the following members:  <pre> <b>struct netbuf</b> <i>addr</i>; <b>struct netbuf</b> <i>opt</i>; <b>struct netbuf</b> <i>udata</i>; </pre> <p>On return from this call:</p> <p><i>addr</i> Specifies the protocol address of the sending user.</p> <p><i>opt</i> Identifies options that were associated with this data unit.</p> <p><i>udata</i> Specifies the user data that was received.</p> <p>The <i>maxlen</i> field of <i>addr</i>, <i>opt</i>, and <i>udata</i> must be set before calling this subroutine to indicate the maximum size of the buffer for each.</p>
<i>flags</i>	Indicates that the complete data unit was not received.

## Valid States

**T\_IDLE**

## Return Values

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBODATA</b>	<b>O_NONBLOCK</b> was set, but no data units are currently available from the transport provider.
<b>TBUFOVFLW</b>	The number of bytes allocated for the incoming protocol address or options ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the information. The unit data information to be returned in the <i>unitdata</i> parameter is discarded.
<b>TLOOK</b>	An asynchronous event has occurred on the transport endpoint and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **fcntl** subroutine, **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvuderr** subroutine, **t\_sndudata** subroutine.

---

## t\_rcvuderr Subroutine for X/Open Transport Interface

### Purpose

Receive a unit data error indication.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_rcvuderr (fd, uderr)
int fd;
struct t_uderr *uderr;
```

### Description

The **t\_rcvuderr** subroutine is used in connectionless mode to receive information concerning an error on a previously sent data unit, and should only be issued following a unit data error indication. It informs the transport user that a data unit with a specific destination address and protocol options produced an error.

### Parameters

*fd*            Identifies the local transport endpoint through which the error report will be received.  
*uderr*        Points to a **t\_uderr** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
long error;
```

The *maxlen* field of *addr* and *opt* must be set before calling this subroutine to indicate the maximum size of the buffer for each.

On return from this call:

*addr*        Specifies the destination protocol address of the erroneous data unit.  
*opt*        Identifies options that were associated with the data unit.  
*error*       Specifies a protocol-dependent error code.

If the user does not care to identify the data unit that produced an error, *uderr* may be set to a null pointer, and the **t\_rcvuderr** subroutine simply clears the error indication without reporting any information to the user.

### Valid States

**T\_IDLE**

### Return Values

0            Successful completion.  
-1        **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBUFOVFLW</b>	The number of bytes allocated for the incoming protocol address or options ( <i>maxlen</i> ) is greater than 0 but not sufficient to store the information. The unit data information to be returned in the <i>uderr</i> parameter is discarded.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TNOUDERR</b>	No unit data error indication currently exists on the specified transport endpoint.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_rcvudata** subroutine, **t\_sndudata** subroutine.

---

## t\_snd Subroutine for X/Open Transport Interface

### Purpose

Send data or expedited data over a connection.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
```

```
int t_snd (  
    int fd,  
    void *buf,  
    unsigned int nbytes,  
    int *flags)
```

### Description

The **t\_snd** subroutine is used to send either normal or expedited data. By default, the **t\_snd** subroutine operates in synchronous mode and may wait if flow control restrictions prevents the data from being accepted by the local transport provider at the time the call is made. However, if **O\_NONBLOCK** is set (via the **t\_open** subroutine or *fcntl*), the **t\_snd** subroutine executes in asynchronous mode, and fails immediately if there are flow control restrictions. The process can arrange to be informed when the flow control restrictions are cleared via either the **t\_look** subroutine or the Event Management interface.

On successful completion, the **t\_snd** subroutine returns the number of bytes accepted by the transport provider. Normally this equals the number of bytes specified in the *nbytes* parameter. However, if **O\_NONBLOCK** is set, it is possible that only part of the data is actually accepted by the transport provider. In this case, the **t\_snd** subroutine returns a value that is less than the value of the *nbytes* parameter. If the value of the *nbytes* parameter is zero and sending of zero octets is not supported by the underlying transport service, the **t\_snd** subroutine returns -1 with **t\_errno** set to **TBADDATA**.

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. Therefore if several processes issue concurrent **t\_snd** calls then the different data may be intermixed.

Multiple sends which exceed the maximum TSDU or ETSDU size may not be discovered by the X/Open Transport Interface. In this case an implementation-dependent error will result (generated by the transport provider) perhaps on a subsequent XTI call. This error may take the form of a connection abort, a **TSYSERR**, a **TBADDATA** or a **TPROTO** error.

If multiple sends which exceed the maximum TSDU or ETSDU size are detected by the X/Open Transport Interface, **t\_snd** fails with **TBADDATA**.

## Parameters

*fd* Identifies the local transport endpoint over which data should be sent.  
*buf* Points to the user data.  
*nbytes* Specifies the number of bytes of user data to be sent.  
*flags* Specifies any optional flags described below:

### **T\_EXPEDITED**

If set in the *flags* parameter, the data is sent as expedited data and is subject to the interpretations of the transport provider.

### **T\_MORE**

If set in the *flags* parameter, indicates to the transport provider that the transport service data unit (TSDU) (or expedited transport service data unit - ETSDU) is being sent through multiple **t\_snd** calls. Each **t\_snd** call with the **T\_MORE** flag set indicates that another **t\_snd** call will follow with more data for the current TSDU (or ETSDU).

The end of the TSDU (or ETSDU) is identified by a **t\_snd** call with the **T\_MORE** flag not set. Use of **T\_MORE** enables a user to break up large logical data units without losing the boundaries of those units at the other end of the connection. The flag implies nothing about how the data is packaged for transfer below the transport interface. If the transport provider does not support the concept of a TSDU, as indicated in the *info* parameter on return from the **t\_open** or **t\_getinfo** subroutines, the **T\_MORE** flag is not meaningful and is ignored if set.

The sending of a zero-length fragment of a TSDU or ETSDU is only permitted where this is used to indicate the end of a TSDU or ETSDU, for example, when the **T\_MORE** flag is not set. Some transport providers also forbid zero-length TSDUs and ETSDUs. See Appendix A, ISO Transport Protocol Information for a fuller explanation.

## Valid States

**T\_DATAXFER**, **T\_INREL**.

## Return Values

On successful completion, the **t\_snd** subroutine returns the number of bytes accepted by the transport provider. Otherwise, -1 is returned on failure and **t\_errno** is set to indicate the error.

Note, that in asynchronous mode, if the number of bytes accepted by the transport provider is less than the number of bytes requested, this may indicate that the transport provider is blocked due to flow control.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.

Value	Description
<b>TBADDATA</b>	Illegal amount of data: <ul style="list-style-type: none"> <li>A single send was attempted specifying a TSDU (ETSDU) or fragment TSDU (ETSDU) greater than that specified by the current values of the TSDU or ETSDU fields in the <i>info</i> argument;</li> <li>a send of a zero byte TSDU (ETSDU) or zero byte fragment of a TSDU (ETSDU) is not supported by the provider (see Appendix A, ISO Transport Protocol Information) .</li> <li>multiple sends were attempted resulting in a TSDU (ETSDU) larger than that specified by the current value of the TSDU or ETSDU fields in the <i>info</i> argument - the ability of an XTI implementation to detect such an error case is implementation-dependent. See "Implementation Specifics".</li> </ul>
<b>TBADFLAG</b>	An invalid flag was specified.
<b>TFLOW</b>	<b>O_NONBLOCK</b> was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine, **t\_rcv** subroutine.

---

## t\_snddis Subroutine for X/Open Transport Interface

### Purpose

Send user-initiated disconnect request.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
```

```
int t_snddis (
    int fd,
    const struct t_call *call)
```

### Description

The **t\_snddis** subroutine is used to initiate an abortive release on an already established connection, or to reject a connect request.

The **t\_snddis** subroutine is an abortive disconnect. Therefore a **t\_snddis** call issued on a connection endpoint may cause data previously sent via the **t\_snd** subroutine, or data not yet received, to be lost (even if an error is returned).

### Parameters

*fd*        Identifies the local transport endpoint of the connection.

*call* Specifies information associated with the abortive release. The *call* parameter points to a **t\_call** structure which contains the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
int sequence;
```

The values in the *call* parameter have different semantics, depending on the context of the call to the **t\_snddis** subroutine. When rejecting a connect request, the *call* parameter must be non-null and contain a valid value of *sequence* to uniquely identify the rejected connect indication to the transport provider. The *sequence* field is only meaningful if the transport connection is in the **T\_INCON** state. The *addr* and *opt* fields of the *call* parameter are ignored. In all other cases, the *call* parameter need only be used when data is being sent with the disconnect request. The *addr*, *opt* and *sequence* fields of the **t\_call** structure are ignored. If the user does not wish to send data to the remote user, the value of the *call* parameter may be a null pointer.

The *udata* field specifies the user data to be sent to the remote user. The amount of user data must not exceed the limits supported by the transport provider, as returned in the the **t\_open** or **t\_getinfo** subroutines *info* parameter *discon* field. If the *len* field of *udata* is zero, no data will be sent to the remote user.

## Valid States

**T\_DATAXFER**, **T\_OUTCON**, **T\_OUTREL**, **T\_INREL**, **T\_INCON**(*ocnt* > 0).

## Return Values

0 Successful completion.  
-1 **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADDATA</b>	The amount of user data specified was not within the bounds allowed by the transport provider.
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TBADSEQ</b>	An invalid sequence number was specified, or a null <i>call</i> pointer was specified, when rejecting a connect request.
<b>TLOOK</b>	An asynchronous event, which requires attention has occurred.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_connect** subroutine, **t\_getinfo** subroutine, **t\_listen** subroutine, **t\_open** subroutine.

---

## t\_sndrel Subroutine for X/Open Transport Interface

### Purpose

Initiate an orderly release.

## Library

X/Open Transport Interface Library (**libxti.a**)

## Syntax

```
#include <xti.h>
int t_sndrel (fd)
int fd;
```

## Description

The **t\_sndrel** subroutine is used to initiate an orderly release of a transport connection and indicates to the transport provider that the transport user has no more data to send.

After calling the **t\_sndrel** subroutine, the user may not send any more data over the connection. However, a user may continue to receive data if an orderly release indication has not been received. This subroutine is an optional service of the transport provider and is only supported if the transport provider returned service type **T\_COTS\_ORD** on the **t\_open** or **t\_getinfo** subroutines.

## Parameter

*fd*      Identifies the local transport endpoint where the connection exists.

## Valid States

**T\_DATAXFER**, **T\_INREL**.

## Return Values

0          Successful completion.  
-1        **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TFLOW</b>	<b>O_NONBLOCK</b> was set, but the flow control mechanism prevented the transport provider from accepting the subroutine at this time.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint and requires immediate attention.
<b>TNOTSUPPORT</b>	This subroutine is not supported by the underlying transport provider.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.

## Related Information

The **t\_getinfo** subroutine, **t\_open** subroutine, **t\_rcvrel** subroutine.

---

## t\_sndudata Subroutine for X/Open Transport Interface

### Purpose

Send a data unit.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

int t_sndudata (
    int fd,
    const struct t_unitdata *unitdata)
```

### Description

The **t\_sndudata** subroutine is used in connectionless mode to send a data unit from another transport user.

By default, the **t\_sndudata** subroutine operates in synchronous mode and waits if flow control restrictions prevents the data from being accepted by the local transport provider at the time the call is made. However, if **O\_NONBLOCK** is set (via the **t\_open** subroutine or *fcntl*), the **t\_sndudata** subroutine executes in asynchronous mode and fails under such conditions. The process can arrange to be notified of the clearance of a flow control restriction via either the **t\_look** subroutine or the Event Management interface.

If the amount of data specified in the *udata* field exceeds the TSDU size as returned in the **t\_open** or **t\_getinfo** subroutines *info* parameter *tsdu* field, a **TBADDATA** error will be generated. If the **t\_sndudata** subroutine is called before the destination user has activated its transport endpoint (see the **t\_bind** subroutine), the data unit may be discarded.

If it is not possible for the transport provider to immediately detect the conditions that cause the errors **TBADDADDR** and **TBADOPT**. These errors will alternatively be returned by the **t\_rcvuderr** subroutine. Therefore, an application must be prepared to receive these errors in both of these ways.

### Parameters

*fd* Identifies the local transport endpoint through which data will be sent.

*unitdata* Points to a **t\_unitdata** structure containing the following members:

```
struct netbuf addr;
struct netbuf opt;
struct netbuf udata;
```

In the *unitdata* structure:

*addr* Specifies the protocol address of the destination user.

*opt* Identifies options that the user wants associated with this request. The user may choose not to specify what protocol options are associated with the transfer by setting the *len* field of *opt* to zero. In this case, the provider may use default options.

*udata* Specifies the user data to be sent. If the *len* field of *udata* is zero, and sending of zero octets is not supported by the underlying transport service, the **t\_sndudata** subroutine returns -1 with **t\_errno** set to **TBADDATA**.

## Valid States

T\_IDLE

## Return Values

0        Successful completion.  
-1        **t\_errno** is set to indicate an error.

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
TBADADDR	The specified protocol address was in an incorrect format or contained illegal information.
TBADDATA	Illegal amount of data. A single send was attempted specifying a TSDU greater than that specified in the <i>info</i> parameter, or a send of a zero byte TSDU is not supported by the provider.
TBADF	The specified file descriptor does not refer to a transport endpoint.
TBADOPT	The specified options were in an incorrect format or contained illegal information.
TFLOW	<b>O_NONBLOCK</b> was set, but the flow control mechanism prevented the transport provider from accepting any data at this time.
TLOOK	An asynchronous event has occurred on the transport endpoint.
TNOTSUPPORT	This subroutine is not supported by the underlying transport provider.
TOUTSTATE	The subroutine was issued in the wrong sequence on the transport endpoint referenced by the <i>fd</i> parameter.
TPROTO	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
TSYSERR	A system error has occurred during execution of this subroutine.

## Related Information

The **fcntl** subroutine, **t\_alloc** subroutine, **t\_open** subroutine, **t\_rcvudata** subroutine, **t\_rcvuderr** subroutine.

---

## t\_strerror Subroutine for X/Open Transport Interface

### Purpose

Produce an error message string.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>

const char *t_strerror (
    int errnum)
```

### Description

The **t\_strerror** subroutine maps the error number to a language-dependent error message string and returns a pointer to the string. The error number specified by the *errnum* parameter corresponds to an X/Open Transport Interface error. The string pointed to is not modified by the program, but may be overwritten by a subsequent call to the **t\_strerror** subroutine. The string is not terminated by a newline.

character. The language for error message strings written by the **t\_strerror** subroutine is implementation-defined. If it is English, the error message string describing the value in **t\_errno** is identical to the comments following the **t\_errno** codes defined in the **xti.h** header file. If an error code is unknown, and the language is English, **t\_strerror** returns the string.

```
"<error>: error unknown"
```

where **<error>** is the error number supplied as input. In other languages, an equivalent text is provided.

## Parameter

*errnum*      Specifies the error number.

## Valid States

ALL - except **T\_UNINIT**.

## Return Values

The **t\_strerror** subroutine returns a pointer to the generated message string.

## Related Information

The **t\_error** subroutine.

---

## t\_sync Subroutine for X/Open Transport Interface

### Purpose

Synchronize transport library.

### Library

X/Open Transport Interface Library (**libxti.a**)

### Syntax

```
#include <xti.h>
int t_sync (fd)
int fd;
```

### Description

The **t\_sync** subroutine synchronizes the data structures managed by the transport library with information from the underlying transport provider. In doing so, if the file descriptor referenced a transport endpoint, the subroutine can convert an uninitialized file descriptor (obtained using the **open** or **dup** subroutines or as a result of a **fork** operation and an **exec** operation) to an initialized transport endpoint, by updating and allocating the necessary library data structures. This subroutine also allows two cooperating processes to synchronize their interaction with a transport provider.

For example, if a process forks a new process and issues an **exec** operation, the new process must issue a **t\_sync** to build the private library data structure associated with a transport endpoint and to synchronize the data structure with the relevant provider information.

It is important to remember that the transport provider treats all users of a transport endpoint as a single user. If multiple processes are using the same endpoint, they should coordinate their activities so as not to violate the state of the transport endpoint. The **t\_sync** subroutine returns the current state of the transport endpoint to the user, thereby enabling the user to verify the state before taking further action. This

coordination is only valid among cooperating processes; it is possible that a process or an incoming event could change the endpoint's state *after* a **t\_sync** call is issued.

If the transport endpoint is undergoing a state transition when the **t\_sync** subroutine is called, the subroutine will fail.

## Parameter

*fd* Specifies the transport endpoint.

## Valid States

ALL - except **T\_UNINIT**.

## Return Values

On successful completion, the state of the transport endpoint is returned. Otherwise, a value of -1 is returned and **t\_errno** is set to indicate an error. The state returned is one of the following:

Value	Description
<b>T_UNBND</b>	Unbound.
<b>T_IDLE</b>	Idle.
<b>T_OUTCON</b>	Outgoing connection pending.
<b>T_INCON</b>	Incoming connection pending.
<b>T_DATAXFER</b>	Data transfer.
<b>T_OUTREL</b>	Outgoing orderly release (waiting for an orderly release indication).
<b>T_INREL</b>	Incoming orderly release (waiting for an orderly release request).

## Error Codes

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint. This error may be returned when the <i>fd</i> parameter has been previously closed or an erroneous number may have been passed to the call.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).
<b>TSTATECHNG</b>	The transport endpoint is undergoing a state change.
<b>TSYSERR</b>	A system error has occurred during execution of this function.

## Related Information

The **dup** subroutine, **exec** subroutine, **fork** subroutine, **open** subroutine.

---

## t\_unbind Subroutine for X/Open Transport Interface

### Purpose

Disable a transport endpoint.

### Library

X/Open Transport Interface Library (**libxti.a**)

## Syntax

```
#include <xti.h>
int t_unbind (fd)
int fd;
```

## Description

The **t\_unbind** subroutine disables the transport endpoint which was previously bound by **t\_bind**. On completion of this call, no further data or events destined for this transport endpoint are accepted by the transport provider. An endpoint which is disabled by using the **t\_unbind** subroutine can be enabled by a subsequent call to the **t\_unbind** subroutine.

## Parameter

*fd* Specifies the transport endpoint.

## Valid States

**T\_IDLE**

## Return Values

0 Successful completion.  
-1 **t\_errno** is set to indicate an error.

## Errors

On failure, **t\_errno** is set to one of the following:

Value	Description
<b>TBADF</b>	The specified file descriptor does not refer to a transport endpoint.
<b>TOUTSTATE</b>	The subroutine was issued in the wrong sequence.
<b>TLOOK</b>	An asynchronous event has occurred on this transport endpoint.
<b>TSYSERR</b>	A system error has occurred during execution of this subroutine.
<b>TPROTO</b>	This error indicates that a communication problem has been detected between the X/Open Transport Interface and the transport provider for which there is no other suitable X/Open Transport Interface ( <b>t_errno</b> ).

## Related Information

The **t\_bind** subroutine.

---

## Options for the X/Open Transport Interface

Options are formatted according to the **t\_opthdr** structure as described in "Use of Options for the X/Open Transport Interface". A transport provider compliant to this specification supports none, all, or any subset of the options defined in the following sections: "TCP/IP-Level Options" to "IP-level Options". An implementation may restrict the use of any of these options by offering them only in the privileged or read-only mode.

## TCP-Level Options

The protocol level is **INET\_TCP**. For this level, the following table shows the options that are defined.

### *TCP-Level Options*

Option Name	Type of Option Value	Legal Option Value	Meaning
TCP_KEEPAIVE	struct t_kpalive	see text following table	check if connections are live
TCP_MAXSEG	unsigned long	length in octets	get TCP maximum segment size
TCP_NODELAY	unsigned long	T_YES T_NO	don't delay send to coalesce packets

### **TCP\_KEEPAIVE**

If set, a keep-alive timer is activated to monitor idle connections that may no longer exist. If a connection has been idle since the last keep-alive timeout, a keep-alive packet is sent to check if the connection is still alive or broken.

Keep-alive packets are not an explicit feature of TCP, and this practice is not universally accepted. According to **RFC 1122**:

"a keep-alive mechanism should only be invoked in server applications that might otherwise hang indefinitely and consume resources unnecessarily if a client crashes or aborts a connection during a network failure."

The option value consists of a structure **t\_kpalive** declared as:

```
struct t_kpalive {
    long kp_onoff;
    long kp_timeout;
}
```

The **t\_kpalive** fields and the possible values are:

#### *kp\_onoff*

Switches option on or off. Legal values for the field are:

**T\_NO** Switch keep-alive timer off.

**T\_YES** Activate keep-alive timer.

#### **T\_YES | T\_GARBAGE**

Activate keep-alive timer and send garbage octet.

Usually, an implementation should send a keep-alive packet with no data (**T\_GARBAGE** not set). If **T\_GARBAGE** is set, the keep-alive packet contains one garbage octet for compatibility with erroneous TCP implementations.

An implementation is, however, not obliged to support **T\_GARBAGE** (see RFC 1122). Since the *kp\_onoff* value is an absolute requirement, the request "**T\_YES | T\_GARBAGE**" may therefore be rejected.

#### *kp\_timeout*

Specifies the keep-alive timeout in minutes. This field determines the frequency of keep-alive packets being sent, in minutes. The transport user can request the default value by setting the field to **T\_UNSPEC**. The default is implementation-dependent, but at least 120 minutes (see RFC 1122). Legal values for this field are **T\_UNSPEC** and all positive numbers.

The timeout value is not an absolute requirement. The implementation may pose upper and lower limits to this value. Requests that fall short of the lower limit may be negotiated to the lower limit.

The use of this option might be restricted to privileged users.

### **TCP\_MAXSEG**

Used to retrieve the maximum TCP segment size. This option is read-only.

## TCP\_NODELAY

Under most circumstances, TCP sends data as soon as it is presented. When outstanding data has not yet been acknowledged, it gathers small amounts of output to be sent in a single packet once an acknowledgment is received. For a small number of clients, such as window systems (for example, Enhanced AIXwindows) that send a stream of mouse events which receive no replies, this packetization may cause significant delays. **TCP\_NODELAY** is used to defeat this algorithm. Legal option values are:

**T\_YES** Do not delay.

**T\_NO** Delay.

These options are not association-related. The options may be negotiated in all X/Open Transport Interface states except **T\_UNBIND** and **T\_UNINIT**. The options are read-only in the **T\_UNBIND** state. See "The Use of Options for the X/Open Transport Interface" for the differences between association-related options and those options that are not.

## Absolute Requirements

A request for **TCP\_NODELAY** and a request to activate **TCP\_KEEPALIVE** is an absolute requirement. **TCP\_MAXSEG** is a read-only option.

## UDP-level Options

The protocol level is **INET\_UDP**. The option defined for this level is shown in the following table.

### UDP-Level Options

Option Name	Type of Option Value	Legal Option Value	Meaning
UDP_CHECKSUM	unsigned long	T_YES/T_NO	checksum computation

## UDP\_CHECKSUM

Allows disabling and enabling of the UDP checksum computation. The legal values are:

**T\_YES** Checksum enabled.

**T\_NO** Checksum disabled.

This option is association-related. It may be negotiated in all XTI states except **T\_UNBIND** and **T\_UNINIT**. It is read-only in state **T\_UNBND**.

If this option is returned with the **t\_rcvudata** subroutine, its value indicates whether a checksum was present in the received datagram or not.

Numerous cases of undetected errors have been reported when applications chose to turn off checksums for efficiency. The advisability of ever turning off the checksum check is very controversial.

## Absolute Requirements

A request for this option is an absolute requirement.

## IP-level Options

The protocol level is **INET\_IP**. The options defined for this level are listed in the following table.

### IP-Level Options

Option Name	Type of Option Value	Legal Option Value	Meaning
IP_BROADCAST	unsigned int	T_YES/T_NO	permit sending of broadcast messages
IP_DONTROUTE	unsigned int	T_YES/T_NO	just use interface addresses
IP_OPTIONS	array of unsigned characters	see text	IP per-packet options
IP_REUSEADDR	unsigned int	T_YES/T_NO	allow local address reuse

### IP-Level Options

Option Name	Type of Option Value	Legal Option Value	Meaning
IP_TOS	unsigned char	see text	IP per-packet type of service
IP_TTL	unsigned char	time in seconds	IP per packet time-to-live

#### IF\_BROADCAST

Requests permission to send broadcast datagrams. It was defined to make sure that broadcasts are not generated by mistake. The use of this option is often restricted to privileged users.

#### IP\_DONTROUTE

Indicates that outgoing messages should bypass the standard routing facilities. It is mainly used for testing and development.

#### IP\_OPTIONS

Sets or retrieves the OPTIONS field of each outgoing (incoming) IP datagram. Its value is a string of octets composed of a number of IP options, whose format matches those defined in the IP specification with one exception: the list of addresses for the source routing options must include the first-hop gateway at the beginning of the list of gateways. The first-hop gateway address will be extracted from the option list and the size adjusted accordingly before use.

The option is disabled if it is specified with "no value," for example, with an option header only.

The **t\_connect** (in synchronous mode), **t\_listen**, **t\_rcvconnect** and **t\_rcvudata** subroutines return the OPTIONS field, if any, of the received IP datagram associated with this call. The **t\_rcvuderr** subroutine returns the OPTIONS field of the data unit previously sent that produced the error. The **t\_optmgmt** subroutine with **T\_CURRENT** set retrieves the currently effective **IP\_OPTIONS** that is sent with outgoing datagrams.

Common applications never need this option. It is mainly used for network debugging and control purposes.

#### IP\_REUSEADDR

Many TCP implementations do not allow the user to bind more than one transport endpoint to addresses with identical port numbers. If **IP\_REUSEADDR** is set to **T\_YES** this restriction is relaxed in the sense that it is now allowed to bind a transport endpoint to an address with a port number and an underspecified internet address ("wild card" address) and further endpoints to addresses with the same port number and (mutually exclusive) fully specified internet addresses.

## IP\_TOS

Sets or retrieves the *type-of-service* field of an outgoing (incoming) IP datagram. This field can be constructed by any OR'ed combination of one of the precedence flags and the type-of-service flags **T\_LDELAY**, **T\_HITHRPT**, and **T\_HIREL**:

- Precedence:

These flags specify datagram precedence, allowing senders to indicate the importance of each datagram. They are intended for Department of Defense applications. Legal flags are:

**T\_ROUTINE**  
**T\_PRIORITY**  
**T\_IMMEDIATE**  
**T\_FLASH**  
**T\_OVERRIDEFLASH**  
**T\_CRITIC\_ECP**  
**T\_INETCONTROL**  
**T\_NETCONTROL**

Applications using **IP\_TOS** but not the precedence level should use the value **T\_ROUTINE** for precedence.

- Type of service:

These flags specify the type of service the IP datagram desires. Legal flags are:

### **T\_NOTOS**

requests no distinguished type of service

### **T\_LDELAY**

requests low delay

### **T\_HITHRPT**

requests high throughput

### **T\_HIREL**

requests high reliability

The option value is set using the macro **SET\_TOS(prec, tos)** where *prec* is set to one of the precedence flags and *tos* to one or an OR'ed combination of the type-of-service flags. **SET\_TOS** returns the option value.

The **t\_connect**, **t\_listen**, **t\_rcvconnect** and **t\_rcvudata** subroutines return the *type-of-service* field of the received IP datagram associated with this call. The **t\_rcvuderr** subroutine returns the *type-of-service* field of the data unit previously sent that produced the error.

The **t\_optmgmt** subroutine with **T\_CURRENT** set retrieves the currently effective **IP\_TOS** value that is sent with outgoing datagrams.

The requested *type-of-service* cannot be guaranteed. It is a hint to the routing algorithm that helps it choose among various paths to a destination. Note also, that most hosts and gateways in the Internet these days ignore the *type-of-service* field.

## IP\_TIL

This option is used to set the *time-to-live* field in an outgoing IP datagram. It specifies how long, in seconds, the datagram is allowed to remain in the Internet. The *time-to-live* field of an incoming datagram is not returned by any function (since it is not an association-related option).

**IP\_OPTIONS** and **IP\_TOS** are both association-related options. All other options are not association-related.

**IP\_REUSEADDR** may be negotiated in all XTI states except **T\_UNINIT**. All other options may be negotiated in all other XTI states except **T\_UNBND** and **T\_UNINIT**; they are read-only in the state **T\_UNBND**.

## Absolute Requirements

A request for any of these options in an absolute requirement.



---

## Chapter 4. Packet Capture

---

### Packet Capture Library Subroutines

The packet capture library contains subroutines that allow users to communicate with the packet capture facility provided by the operating system to read unprocessed network traffic. Applications using these subroutines must be run as root. These subroutines are maintained in the **libpcap.a** library:

- pcap\_close
- pcap\_compile
- pcap\_datalink
- pcap\_dispatch
- pcap\_dump
- pcap\_dump\_close
- pcap\_dump\_open
- pcap\_file
- pcap\_fileno
- pcap\_geterr
- pcap\_is\_swapped
- pcap\_lookupdev
- pcap\_lookupnet
- pcap\_loop
- pcap\_major\_version
- pcap\_minor\_version
- pcap\_next
- pcap\_open\_live
- pcap\_open\_offline
- pcap\_perror
- pcap\_setfilter
- pcap\_snapshot
- pcap\_stats
- pcap\_strerror

“ioctl BPF Control Operations” performs packet-capture-related control operations.

---

### ioctl BPF Control Operations

#### Purpose

Performs packet-capture-related control operations.

#### Syntax

```
#include <sys/ioctl.h>
int ioctl ( int fd, int cmd [, arg ] )
```

#### Description

The Berkeley Packet Filter (BPF) ioctl commands perform a variety of packet-capture-related control. The *fd* argument is a BPF device descriptor. For non-packet-capture descriptors, functions performed by this call are unspecified.

The *cmd* parameter and an optional third parameter (with varying types) are passed to and interpreted by the BPF ioctl function to perform an appropriate control operation that is specified by the user.

## Parameters

<i>fd</i>	Specifies an open file descriptor that refers to a BPF device created using the open call.
<i>cmd</i>	Selects the control function to be performed.
<i>arg</i>	Represents additional information that is needed to perform the requested function. The type of the <i>arg</i> parameter is either an integer or a pointer to a BPF-specific data structure, depending on the particular control request.

## BPF Control Operations

In addition to the **FIONREAD** ioctl command, the following commands can be applied to any open BPF device. The *arg* parameter is a pointer to the indicated type.

ioctl command	Type of the <i>arg</i> parameter	Description
<b>BIOCGBLLEN</b>	<b>u_int</b>	Returns the buffer length for reads on BPF devices.
<b>BIOCSBLEN</b>	<b>u_int</b>	Sets the buffer length for reads on BPF devices. The <i>buffer</i> parameter must be set before the device is attached to an interface with the <b>BIOCSETIF</b> command. If the requested buffer size cannot be accommodated, the closest allowable size is set and returned in the <i>arg</i> parameter.
<b>BIOCGDLT</b>	<b>u_int</b>	Returns the type of the data link layer underlying the attached interface.
<b>BIOCPROMISC</b>	N/A	Forces the interface into promiscuous mode. All packets, not just those destined for the local host, are processed. A listener that opened its interface nonpromiscuously can receive packets promiscuously, because more than one device can be listening on a given interface. The problem can be remedied with an appropriate filter.
<b>BIOCFLUSH</b>	N/A	Flushes the buffer of incoming packets, and resets the statistics that are returned by the <b>BIOCGSTATS</b> command.
<b>BIOCGETIF</b>	<b>struct ifreq</b>	Returns the name of the hardware interface that the device is listening on. The name is returned in the <b>ifr_name</b> field of the <b>ifreq</b> structure. All other fields are undefined.
<b>BIOCSETIF</b>	<b>struct ifreq</b>	Sets the hardware interface associate with the device. This command must be performed before any pack-packets can be read. The device is indicated by the name using the <b>ifr_name</b> field of the <b>ifreq</b> structure. Additionally, the command performs the actions of the <b>BIOCFLUSH</b> command.
<b>BIOCGRTIMEOUT</b>	<b>struct timeval</b>	Gets the read timeout value. The <i>arg</i> parameter specifies the length of time to wait before a read request times out. This parameter is initialized to zero by an open, indicating no timeout.
<b>BIOCSRTIMEOUT</b>	<b>struct timeval</b>	Sets the read timeout value described in the <b>BIOCGRTIMEOUT</b> command.
<b>BIOCGSTATS</b>	<b>struct bpf_stat</b>	Returns the a structure of packet statistics. The structure is defined in the <b>net/bpf.h</b> file.
<b>BIOCIMMEDIATE</b>	<b>u_int</b>	Enables or disables the <b>immediate</b> mode, based on the truth value of the <i>arg</i> parameter. When the <b>immediate</b> mode is enabled, reads return immediately upon packet reception. Otherwise, a read will be blocked until either the kernel buffer becomes full or a timeout occurs.

ioctl command	Type of the <i>arg</i> parameter	Description
<b>BIOCSETF</b>	<b>struct bpf_program</b>	Sets the filter program used by the kernel to discard uninteresting packets. The <b>bpf_program</b> structure is defined in the <b>net/bpf.h</b> file.
<b>BIOCVERSION</b>	<b>struct bpf_version</b>	Returns the major and minor version numbers of the filter language currently recognized by the kernel. Before installing a filter, applications must check that the current version is compatible with the running kernel. The current version numbers are given by the <b>BPF_MAJOR_VERSION</b> and <b>BPF_MINOR_VERSION</b> variables from the <b>net/bpf.h</b> file. An incompatible filter might result in undefined behavior.

## Return Values

Upon successful completion, `ioctl` returns a value of 0. Otherwise, it returns a value of -1 and sets **errno** to indicate the error.

## Error Codes

The `ioctl` commands fail under the following general conditions:

<b>EINVAL</b>	A command or argument, which is not valid, was specified.
<b>ENETDOWN</b>	The underlying interface or network is down.
<b>ENXIO</b>	The underlying interface is not found.
<b>ENOBUFS</b>	Insufficient memory was available to process the request.
<b>EEXIST</b>	The BPF device already exists.
<b>ENODEV</b>	The BPF device could not be set up.
<b>EINTR</b>	A signal was caught during an <code>ioctl</code> operation.
<b>EACCES</b>	The permission was denied for the specified operation.
<b>EADDRNOTAVAIL</b>	The specified address is not available for interface.
<b>ENOMEM</b>	The available memory is not enough.
<b>ESRCH</b>	Such a process does not exist.

## Related Information

The Packet Capture Library Overview in the *AIX Version 6.1 Communications Programming Concepts*.



---

## Appendix. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Dept. LRAS/Bldg. 003  
11400 Burnet Road  
Austin, TX 78758-3498  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be the trademarks or service marks of others.

---

# Index

## Special characters

- `_getlong` subroutine 25
- `_getshort` subroutine 26
- `_ll_log` subroutine 3
- `_putlong` subroutine 27
- `_putshort` subroutine 28
- `/etc/hosts` file
  - closing 45, 46
  - opening 206, 207
  - retrieving host entries 74, 75, 76, 78, 80
  - setting file markers 206, 207
- `/etc/networks` file
  - closing 47, 48
  - opening 210, 211
  - retrieving network entries 87, 88, 89, 90, 91, 92
  - setting file markers 210, 211
- `/etc/protocols` file
  - closing 49, 50
  - opening 212, 213
  - retrieving protocol entries 95, 96, 97, 98, 99, 100
  - setting file markers 212, 213
- `/etc/resolv.conf` file
  - retrieving host entries 74, 75, 76, 78
  - searching for domain names 173
  - searching for Internet addresses 173
- `/etc/services` file
  - closing 50, 51
  - opening 108, 109, 214, 215
  - reading 108, 109
  - retrieving service entries 103, 104, 105, 106
  - setting file markers 214, 215

## A

- `accept` subroutine 29
- `adjmsg` utility 261
- administrative operations
  - providing interface for 330
- `allocb` utility 261, 263
- `arp` subroutines
  - `arpresolve_common` 30
  - `arpupdate` 31
- `arpresolve_common` subroutine 30
- `arpupdate` subroutine 31
- ASCII strings
  - converting to Internet addresses 143
- asynchronous mode
  - sending data 373

## B

- `backq` utility 262
- `bcanput` utility 263
- `bind` subroutine 33
- `bufcall` utility 263, 386
- byte streams
  - placing long byte quantities 27

- byte streams (*continued*)
  - placing short byte quantities 28

## C

- `canput` utility 265
- clients
  - server authentication 188
- clone device driver 265
- code, terminating section 333
- communications kernel service subroutines
  - `res_ninit` 176
- compressed domain names
  - expanding 40
- `connect` subroutine 35
- connected sockets
  - creating pairs 243
  - receiving messages 167
  - sending messages 193, 196
- connection requests
  - accepting 343
  - listening 358
  - receiving confirmation 366
- connectionless mode
  - receiving data 371
  - receiving error data 372
  - sending data 377
- converter subroutines
  - `inet_net_ntop` 138
  - `inet_net_pton` 139
  - `inet_ntop` 144
  - `inet_pton` 145
- `copyb` utility 266
- `copymsg` utility 266
- `CreateIoCompletionPort` Subroutine 37
- current domain names
  - returning 73
  - setting 205
- current host identifiers
  - retrieving 81

## D

- data
  - receiving normal or expedited 365
  - sending over connection 373
- data blocks
  - allocating 271
- data link provider, providing interface 268
- `datamsg` utility 267
- default domains
  - searching names 172
- disconnects
  - identifying cause and retrieving data 368
  - user-initiated requests 375
- `dlpi STREAMS` driver 268
- `dn_comp` subroutine 38

- dn\_expand subroutine 40
- domain names
  - compressing 38
- drivers
  - installing 335
  - setting processor levels 333
- dupb utility 269
- dupmsg utility 269

## E

- enableok utility 270
- endhostent subroutine 45
- endhostent\_r subroutine 46
- endnetent subroutine 47
- endnetent\_r subroutine 47
- endnetgrent subroutine 146
- endnetgrent\_r subroutine 48
- endprotoent subroutine 49
- endprotoent\_r subroutine 50
- endservent subroutine 50
- endservent\_r subroutine 51
- enrecvfrom subroutine 52
- enrecvmsg subroutine 52
- erecv subroutine 52
- erecvfrom subroutine 52
- erecvmsg subroutine 52
- error logs
  - generating messages 341
- error messages
  - producing 352
- esballoc utility 271
- esend subroutine 54
- esendmsg subroutine 54
- esendto subroutine 54
- ether\_aton subroutine 56
- ether\_hostton subroutine 56
- ether\_line subroutine 56
- ether\_ntoa subroutine 56
- ether\_ntohost subroutine 56
- event traces
  - generating messages 341

## F

- file descriptors
  - testing 309
- flow control
  - testing priority band 263
- flushband utility 272
- flushq utility 272
- freeaddrinfo subroutine 69
- freeb utility 273
- freemsg utility 274
- functions
  - scheduling calls 382

## G

- get\_auth\_method subroutine
  - authentication methods
    - list of 72
- getaddrinfo subroutine 69
- getadmin utility 274
- getdomainname subroutine 73
- gethostbyaddr subroutine 74
- gethostbyaddr\_r subroutine 75
- gethostbyname subroutine 76
- gethostbyname\_r subroutine 78
- gethostent subroutine 80
- gethostent\_r subroutine 80
- gethostid subroutine 81
- gethostname subroutine 82
- getmid utility 275
- getmsg system call 275
- GetMultipleCompletionStatus Subroutine 83
- getnameinfo subroutine 86
- getnetbyaddr subroutine 87
- getnetbyaddr\_r subroutine 88
- getnetbyname subroutine 89
- getnetbyname\_r subroutine 90
- getnetent subroutine 91
- getnetent\_r subroutine 92
- getnetgrent subroutine 146
- getnetgrent\_r subroutine 93
- getpeername subroutine 94
- getpmsg system call 278
- getprotobyname subroutine 95
- getprotobyname\_r subroutine 96
- getprotobynumber subroutine 97
- getprotobynumber\_r subroutine 98
- getprotoent subroutine 99
- getprotoent\_r subroutine 100
- getq utility 279
- GetQueuedCompletionStatus Subroutine 101
- getservbyname subroutine 103
- getservbyname\_r subroutine 104
- getservbyport subroutine 105
- getservbyport\_r subroutine 106
- getservent subroutine 108
- getservent\_r subroutine 109
- getsmuxEntrybyidentity subroutine 1
- getsmuxEntrybyname subroutine 1
- getsockname subroutine 109
- getsockopt subroutine 111
- group network
  - entries in the
    - handling 93, 146, 211

## H

- host machines
  - setting names 209
  - setting unique identifiers 208
- htonl subroutine 119
- htonll subroutine 119
- htons subroutine 120

## I

- I\_ATMARK operation 292
- I\_CANPUT operation 292
- I\_CKBAND operation 293
- I\_FDINSERT operation 293
- I\_FIND operation 294
- I\_FLUSH operation 295
- I\_FLUSHBAND operation 295
- I\_GETBAND operation 296
- I\_GETCLTIME operation 296
- I\_GETSIG operation 297
- I\_GRDOPT operation 297
- I\_LINK operation 297
- I\_LIST operation 298
- I\_LOOK operation 299
- I\_NREAD operation 299
- I\_PEEK operation 300
- I\_PLINK operation 300
- I\_POP operation 301
- I\_PUNLINK operation 302
- I\_PUSH operation 302
- I\_RECVFD operation 303
- I\_SENDFD operation 304
- I\_SETCLTIME operation 304
- I\_SETSIG operation 305
- I\_SRDOPT operation 306
- I\_STR operation 307
- I\_UNLINK operation 308
- I/O Completion Port (IOCP) Kernel Extension
  - CreateCompletionPort 37
  - GetMultipleCompletionStatus 83
  - GetQueuedCompletionStatus 101
  - PostQueuedCompletionStatus 161
  - ReadFile 165
  - WriteFile 258
- if\_freenameindex subroutine 121
- if\_indextoname subroutine 122
- if\_nameindex subroutine 123
- if\_nametoindex subroutine 123
- incoming connections
  - limiting backlog 157
- inet\_addr subroutine 133
- inet\_lnaof subroutine 136
- inet\_makeaddr subroutine 137
- inet\_net\_ntop subroutine 138
- inet\_net\_pton subroutine 139
- inet\_netof subroutine 140
- inet\_network subroutine 141
- inet\_ntoa subroutine 143
- inet\_ntop subroutine 144
- inet\_pton subroutine 145
- inet6\_opt\_append Subroutine 124
- inet6\_opt\_find Subroutine 125
- inet6\_opt\_finish Subroutine 126
- inet6\_opt\_get\_val Subroutine 126
- inet6\_opt\_init Subroutine 127
- inet6\_opt\_next Subroutine 128
- inet6\_opt\_set\_val Subroutine 129
- inet6\_rth\_add Subroutine 129
- inet6\_rth\_getaddr Subroutine 130
- inet6\_rth\_init Subroutine 130

- inet6\_rth\_reverse Subroutine 131
- inet6\_rth\_segments Subroutine 132
- inet6\_rth\_space Subroutine 133
- initializing logging facility variables 2
- initiating SMUX peers 15
- innetr subroutine 146
- insq utility 280
- Internet addresses
  - constructing 137
  - converting 133
  - converting to ASCII strings 143
  - returning network addresses 136
  - searching 172
- Internet numbers
  - converting Internet addresses 133
  - converting network addresses 141
- ioctl commands 147
- ioctl socket control operations 147
- isastream function 309
- isinet\_addr Subroutine 154
- ISODE library
  - extending base subroutines 9
  - initializing logging facility variables 2
  - logging subroutines 3
- isodetailor subroutine 2

## K

- kvalid\_user subroutine
  - DCE principal mapping 156

## L

- library structures
  - allocating 345
  - freeing 353
- linkb utility 309
- listen subroutine 157
- ll\_dbinit subroutine 3
- ll\_hdinit subroutine 3
- ll\_log subroutine 3
- local host names
  - retrieving 82
- long byte quantities
  - retrieving 25
- long integers, converting
  - from host byte order 119
  - from network byte order 158, 159
  - to host byte order 158, 159
  - to network byte order 119

## M

- Management Information Base (MIB)
  - registering a section 16
- mapping
  - Ethernet number 56
- memory management subroutines
  - getaddrinfo 69
  - getnameinfo 86
  - if\_freenameindex 121

- mi\_bufcall Utility 310
- mi\_close\_comm Utility 311
- mi\_next\_ptr Utility 312
- mi\_open\_comm Utility 313
- MIB variables
  - encoding values from 5
  - setting variable values 11
- minor devices, opening on another driver 265
- modules
  - comparing names 294
  - installing 335
  - listing all names on stream 298
  - pushing to top 302
  - removing below stream head 301
  - retrieving name below stream head 299
  - retrieving pointer to write queue 393
  - returning IDs 275
  - returning pointer to 274
  - returning pointer to read queue 328
  - setting processor level 333
  - testing flow control 263
- msgdsiz utility 314
- multiplexed streams
  - connecting 297, 300
  - disconnecting 302, 308

## N

- name servers
  - creating packets 173
  - creating query messages 173
  - retrieving responses 182
  - sending queries 182
- name2inst subroutine 22
- names
  - binding to sockets 33
- network addresses
  - converting 141
  - returning 136
  - returning network numbers 140
- network entries
  - retrieving 91, 92
  - retrieving by address 87, 88
  - retrieving by name 89, 90
- network host entries
  - retrieving 80
  - retrieving by address 74, 75
  - retrieving by name 76, 78
- network host files
  - opening 206, 207
- network services library
  - supporting transport interface functions 385
- next2inst subroutine 22
- nexttot2inst subroutine 22
- noenable utility 270, 315
- ntohl subroutine 158
- ntohll subroutine 159
- ntohs subroutine 160

## O

- o\_subroutines 5
- o\_generic subroutine 5
- o\_igeneric subroutine 5
- o\_integer subroutine 5
- o\_ipaddr subroutine 5
- o\_number subroutine 5
- o\_specific subroutine 5
- o\_string subroutine 5
- object identifier data structure 7
- object tree (OT)
  - freeing 14
  - MIB list 14
- ode2oid subroutine 7
- OID
  - adjusting the values of entries 9
  - converting text strings to 23
  - extending number of entries in 9
  - manipulating entries 9
- OID (object identifier data structure)
  - manipulating the 7
- oid\_cmp subroutine 7
- oid\_cpy subroutine 7
- oid\_extend subroutine 9
- oid\_free subroutine 7
- oid\_normalize subroutine 9
- oid2ode subroutine 7
- oid2ode\_aux subroutine 7
- oid2prim subroutine 7
- Options 443
- OTHERQ utility 315

## P

- peer entries 1
- peer socket names
  - retrieving 94
- pfmod Packet Filter Module
  - upstream data messages, removing 316
- PostQueuedCompletionStatus Subroutine 161
- prim2oid 7
- priority bands
  - checking write status 292
  - flushing messages 272
- processor levels, setting 333
- protocol data unit (PDU) 12
  - sending 17
  - sending an open 18
- protocol entries
  - retrieving 99, 100
  - retrieving by name 95, 96
  - retrieving by number 97, 98
- psap.h file 8
- pullupmsg utility 319
- putbq utility 320
- putctl utility 321
- putctl1 utility 320
- putmsg system call 322
- putnext utility 324
- putpmsg system call 324

putq utility 325

## Q

qenable utility 327

qreply utility 327

qsize utility 328

queries

    awaiting response 180

querying 111

queue bands

    flushing messages 295

## R

rcmd subroutine 162

rcmd\_af Subroutine 164

RD utility 328

read mode

    returning current settings 297

    setting 306

ReadFile Subroutine 165

readobjects subroutine 10

recv subroutine 167

recvfrom subroutine 169

recvmsg subroutine 171

register I/O points

    wantio utility 389

release indications, acknowledging 370

remote hosts

    executing commands 162

    starting command execution 183

reporting errors to log files 3

res\_init subroutine 172

res\_mkquery subroutine 173

res\_ninit subroutine 176

res\_query subroutine 178

res\_search subroutine 180

res\_send subroutine 182

retrieving variables 22

rexec subroutine 183

rexec\_af Subroutine 185

rmvb utility 329

rmvq utility 329

rresvport subroutine 186

rresvport\_af Subroutine 187

ruserok subroutine 188

## S

s\_generic subroutine 11

sad device driver 330

SCTP subroutines

    sctp\_opt\_info 190

    sctp\_peeloff 191

    sctpctrl 192

sctp\_opt\_info subroutine 190

sctp\_peeloff subroutine 191

sctpctrl subroutine 192

send subroutine 193

send\_file

    send the contents of file through a socket 199

send\_file subroutine

    socket options 199

sendmsg subroutine 195

sendto subroutine 197

server query mechanisms

    providing interfaces to 178

service entries

    retrieving by name 103, 104

    retrieving by port 105, 106

service file entries

    retrieving 108, 109

set\_auth\_method subroutine

    authentication methods

        list of 204

setdomainname subroutine 205

sethostent subroutine 206

sethostent\_r subroutine 207

sethostid subroutine 208

sethostname subroutine 209

setnetent subroutine 210

setnetent\_r subroutine 211

setnetgrent subroutine 146

setnetgrent\_r subroutine 211

setprotoent subroutine 212

setprotoent\_r subroutine 213

setservent subroutine 214

setservent\_r subroutine 215

setsockopt subroutine 216

short byte quantities

    retrieving 26

short integers, converting

    from host byte order 120

    from network byte order 160

    to host byte order 160

    to network byte order 120

shutdown subroutine 226

SIGPOLL signal

    informing stream head to issue 305

    returning events of calling process 297

SLP subroutines

    SLPAttrCallback 227

    SLPClose 228

    SLPEscape 229

    SLPFindAttrs 230

    SLPFindScopes 231

    SLPFindSrvs 232

    SLPFindSrvTypes 233

    SLPFree 234

    SLPGetProperty 234

    SLPOpen 235

    SLPParseSrvURL 236

    SLPSrvTypeCallback 239

    SLPSrvURLCallback 240

    SLPUnescape 240

SLPAttrCallback subroutine 227

SLPClose subroutine 228

SLPDereg subroutine 228

SLPEscape subroutine 229

SLPFindAttrs subroutine 230

- SLPFindScopes subroutine 231
- SLPFindSrvs subroutine 232
- SLPFindSrvTypes subroutine 233
- SLPFree subroutine 234
- SLPGetProperty subroutine 234
- SLPOpen subroutine 235
- SLPParseSrvURL subroutine 236
- SLPReg subroutine 237
- SLPRegReport callback subroutine 238
- SLPSrvTypeCallback subroutine 239
- SLPSrvURLCallback subroutine 240
- SLPUnescape subroutine 240
- SMUX
  - communicating with the SNMP agent 17
  - communicating with the snmpd daemon 15
  - ending SNMP communications 12
  - initiating transmission control protocol (TCP) 15
  - peer responsibility level 16
  - reading a MIB variable structure into 10
  - reading the smux\_errno variable 13
  - registering an MIB tree for 16
  - retrieving peer entries 1
  - sending an open PDU 18
  - sending traps to SNMP 20
  - setting debug level for subroutines 15
  - unregistered trees 14
  - waiting for a message 21
- smux\_close subroutine 12
- smux\_error subroutine 13
- smux\_free\_tree subroutine 14
- smux\_init subroutine 15
- smux\_register subroutine 16
- smux\_response subroutine 17
- smux\_simple\_open subroutine 18
- smux\_trap subroutine 20
- smux\_wait subroutine 21
- smux.h file 13
- SNMP multiplexing peers 1
- snmpd daemon
  - incoming messages alert 18
- snmpd.peers file 1
- socket connections
  - accepting 29
  - listening 157
- socket names
  - retrieving 109
- socket options
  - setting 216
- socket receive operations
  - disabling 226
- socket send operations
  - disabling 226
- socket subroutine 241
- socket subroutines
  - freeaddrinfo subroutine 69
  - if\_indextoname subroutine 122
  - if\_nameindex subroutine 123
  - if\_nametoindex subroutine 123
  - inet6\_opt\_append 124
  - inet6\_opt\_find 125
  - inet6\_opt\_finish 126
- socket subroutines (*continued*)
  - inet6\_opt\_get\_val 126
  - inet6\_opt\_init 127
  - inet6\_opt\_next 128
  - inet6\_opt\_set\_val 129
  - inet6\_rth\_add 129
  - inet6\_rth\_getaddr 130
  - inet6\_rth\_init 130
  - inet6\_rth\_reverse 131
  - inet6\_rth\_segments 132
  - inet6\_rth\_space 133
  - rcmd\_af 164
  - rexec\_af 185
  - rresvport\_af 187
- socketpair subroutine 243
- sockets
  - connecting 35
  - creating 241
  - initiating TCP for SMUX peers 15
  - managing 258
  - retrieving with privileged addresses 186
- sockets kernel service subroutines
  - accept 29
  - bind 33
  - connect 35
  - dn\_comp 38
  - getdomainname 73
  - gethostid 81
  - gethostname 82
  - getpeername 94
  - getsockname 109
  - getsockopt 111
  - listen 157
  - recv 167
  - recvfrom 169
  - recvmsg 171
  - send 193
  - sendmsg 195
  - sendto 197
  - setdomainname 205
  - sethostid 208
  - sethostname 209
  - setsockopt 216
  - shutdown 226
  - socket 241
  - socketpair 243
- sockets messages
  - receiving from connected sockets 167
  - receiving from sockets 169, 171
  - sending through any socket 196
- sockets network library subroutines
  - \_getlong 25
  - \_getshort 26
  - \_putlong 27
  - \_putshort 28
  - dn\_expand 40
  - endhostent 45
  - endhostent\_r 46
  - endnetent 47
  - endnetent\_r 47
  - endnetgrent\_r 48

sockets network library subroutines (*continued*)

- endprotoent 49
- endprotoent\_r 50
- endservent 50
- endservent\_r 51
- gethostbyaddr 74
- gethostbyaddr\_r 75
- gethostbyname 76
- gethostbyname\_r 78
- gethostent 80
- gethostent\_r 80
- getnetbyaddr 87
- getnetbyaddr\_r 88
- getnetbyname 89
- getnetbyname\_r 90
- getnetent 91
- getnetent\_r 92
- getprotobyname 95
- getprotobyname\_r 96
- getprotobynumber 97
- getprotobynumber\_r 98
- getprotoent 99
- getprotoent\_r 100
- getservbyname 103
- getservbyname\_r 104
- getservbyport 105
- getservbyport\_r 106
- getservent 108
- getservent\_r 109
- htonl 119
- htonll 119
- htons 120
- inet\_addr 133
- inet\_lnaof 136
- inet\_makeaddr 137
- inet\_netof 140
- inet\_network 141
- inet\_ntoa 143
- ntohl 158
- ntohll 159
- ntohs 160
- rcmd 162
- res\_init 172
- res\_mkquery 173
- res\_query 178
- res\_search 180
- res\_send 182
- rexec 183
- rresvport 186
- ruserok 188
- sethostent 206
- sethostent\_r 207
- setnetent 210
- setnetent\_r 211
- setprotoent 212
- setprotoent\_r 213
- setservent 214
- setservent\_r 215
- sockets-based protocols, providing access 393
- splice subroutine 258
- splstr utility 333

- splx utility 333
- sprintoid subroutine 7
- srv utility 333
  - messages queued 333
- str\_install utility 335
- str2oid subroutine 7
- stream heads
  - checking queue for message 293
  - counting data bytes in first message 299
  - issuing SIGPOLL signal 305
  - removing modules 301
  - retrieving messages 300
  - retrieving module names 299
  - returning set delay time 296
  - setting delay 304
- streamio operations
  - I\_ATMARK 292
  - I\_CANPUT 292
  - I\_CKBAND 293
  - I\_FDINSERT 293
  - I\_FIND 294
  - I\_FLUSH 295
  - I\_FLUSHBAND 295
  - I\_GETBAND 296
  - I\_GETCLTIME 296
  - I\_GETSIG 297
  - I\_GRDOPT 297
  - I\_LINK 297
  - I\_LIST 298
  - I\_LOOK 299
  - I\_NREAD 299
  - I\_PEEK 300
  - I\_PLINK 300
  - I\_POP 301
  - I\_PUNLINK 302
  - I\_PUSH 302
  - I\_RECVFD 303
  - I\_SENDFD 304
  - I\_SETCLTIME 304
  - I\_SETSIG 305
  - I\_SRDOPT 306
  - I\_STR 307
  - I\_UNLINK 308
- STREAMS
  - mi\_bufcall Utility 310
  - mi\_close\_comm Utility 311
  - mi\_next\_ptr Utility 312
  - mi\_open\_comm Utility 313
  - performing control functions 339
  - unweldq Utility 388
  - weldq Utility 392
- STREAMS buffers
  - checking availability 381
- STREAMS device drivers
  - clone 265
  - sad 330
- STREAMS drivers
  - dlpi 268
  - xtiso 393
- STREAMS message blocks
  - copying 266

### STREAMS message blocks *(continued)*

- duplicating descriptors 269
- freeing 273, 274
- removing from head of message 386
- removing from messages 329

### STREAMS messages

- allocating 271
- allocating data blocks 261
- checking buffer availability 381
- checking markings 292
- concatenating 309
- concatenating and aligning data bytes 319
- constructing internal ioctl 307
- converting streamio operations 383
- counting data bytes 299
- creating control 320, 321
- creating, adding information, and sending downstream 293
- determining whether data message 267
- duplicating 269
- flushing in given priority band 272
- generating error-logging and event-tracing 341
- getting next from queue 279
- getting next priority 278
- getting off stream 275
- passing to next queue 324
- placing in queue 280
- putting on queue 325
- removing from queue 329
- retrieving file descriptors 303
- retrieving without removing 300
- returning number of data bytes 314
- returning number on queue 328
- returning priority band of first on queue 296
- returning to beginning of queue 320
- sending 322
- sending in reverse direction 327
- sending priority 324
- sending to stream head at other end of stream pipe 304
- trimming bytes 261

### STREAMS modules

- timod 383
- tirdwr 385

### STREAMS queues

- checking for messages 293
- counting data bytes in first message 299
- enabling 327
- flushing 272
- flushing input or output 295
- getting next message 279
- obtaining information 342
- passing message to next 324
- preventing scheduling 315
- putting messages on 325
- retrieving pointer to write queue 393
- returning message to beginning 320
- returning number of messages 328
- returning pointer to mate 315
- returning pointer to preceding 262
- returning pointer to read queue 328

### STREAMS queues *(continued)*

- returning priority band of first message 296
- scheduling for service 270
- testing for space 265

### STREAMS subroutines

- isastream 309
- t\_accept 343
- t\_alloc 345
- t\_bind 347
- t\_close 349
- t\_connect 350
- t\_error 352
- t\_free 353
- t\_getinfo 354
- t\_getstate 357
- t\_listen 358
- t\_look 360
- t\_open 361
- t\_optmngmt 363
- t\_rcv 365
- t\_rcvconnect 366
- t\_rcvdis 368
- t\_rcvrel 370
- t\_rcvudata 371
- t\_rcvuderr 372
- t\_snd 373
- t\_snddis 375
- t\_sndrel 376
- t\_sndudata 377
- t\_sync 379
- t\_unbind 380

### STREAMS system calls

- getmsg 275
- getpmsg 278
- putmsg 322
- putpmsg 324

### STREAMS utilities

- adjmsg 261
- allcob 261
- backq 262
- bcanput 263
- bufcall 263
- canput 265
- copyb 266
- copymsg 266
- datamsg 267
- dupb 269
- dupmsg 269
- enableok 270
- esballoc 271
- flushband 272
- flushq 272
- freeb 273
- freemsg 274
- getadmin 274
- getmid 275
- getq 279
- insq 280
- linkb 309
- msgdsize 314
- noenable 315

## STREAMS utilities *(continued)*

- OTHERQ 315
- pullupmsg 319
- putbq 320
- putctl 321
- putctl1 320
- putnext 324
- putq 325
- qenable 327
- qreply 327
- qsize 328
- RD 328
- rmvb 329
- rmvq 329
- splstr 333
- splx 333
- str\_install 335
- strlog 341
- strqget 342
- testb 381
- timeout 382
- unbufcall 386
- unlinkb 386
- untimeout 387
- WR 393
- string conversions 23
- strlog utility 341
- strqget utility 342
- synchronous mode
  - sending data 373

## T

- t\_accept subroutine 343
- t\_accept Subroutine 395
- t\_alloc subroutine 345
- t\_alloc Subroutine 397
- t\_bind subroutine 347
- t\_bind Subroutine 399
- t\_close subroutine 349
- t\_close Subroutine 402
- t\_connect subroutine 350
- t\_connect Subroutine 403
- t\_error subroutine 352
- t\_error Subroutine 405
- t\_free subroutine 353
- t\_free Subroutine 406
- t\_getinfo subroutine 354
- t\_getinfo Subroutine 408
- t\_getprotaddr Subroutine 410
- t\_getstate subroutine 357
- t\_getstate Subroutine 411
- t\_listen subroutine 358
- t\_listen Subroutine 412
- t\_look subroutine 360
- t\_look Subroutine 414
- t\_open subroutine 361
- t\_open Subroutine 415
- t\_opthdr 443
- t\_optmgmt subroutine 363
- t\_optmgmt Subroutine 419

- t\_rcv subroutine 365
- t\_rcv Subroutine 425
- t\_rcvconnect subroutine 366
- t\_rcvconnect Subroutine 427
- t\_rcvdis subroutine 368
- t\_rcvdis Subroutine 429
- t\_rcvrel subroutine 370
- t\_rcvrel Subroutine 430
- t\_rcvudata subroutine 371
- t\_rcvudata Subroutine 431
- t\_rcvuderr Subroutine 433
- t\_rdvuderr subroutine 372
- t\_snd subroutine 373
- t\_snd Subroutine 434
- t\_snddis subroutine 375
- t\_snddis Subroutine 436
- t\_sndrel subroutine 376
- t\_sndrel Subroutine 437
- t\_sndudata subroutine 377
- t\_sndudata Subroutine 439
- t\_strerror Subroutine 440
- t\_sync subroutine 379
- t\_sync Subroutine 441
- t\_unbind subroutine 380
- t\_unbind Subroutine 442
- testb utility 381
- text2inst subroutine 22
- text2obj subroutine 23
- text2oid subroutine 23
- timeout utility 382, 387
- timod module 383
- tirdwr module 385
- transport connections, initiating release 376
- transport endpoints
  - binding addresses 347
  - closing 349
  - disabling 380
  - establishing 361
  - establishing connection 350
  - examining current events 360
  - getting current states 357
  - managing options 363
- transport interfaces
  - converting streamio operations into messages 383
  - supporting network services library functions 385
- transport library, synchronizing data 379
- transport protocols
  - getting service information 354
- traps 20

## U

- unbufcall utility 386
- unconnected sockets
  - receiving messages 169
  - sending messages 196, 197
- unique identifiers
  - retrieving 81
- unlinkb utility 386
- untimeout utility 387
- unweldq Utility 388

## **V**

variable bindings 5  
variable initialization 2

## **W**

wantio utility 389  
wantmsg Utility 390  
weldq Utility 392  
WR utility 393  
write queue  
    retrieve a pointer to 393  
WriteFile Subroutine 258

## **X**

xtiso STREAMS driver 393

---

# Readers' Comments — We'd Like to Hear from You

**AIX Version 6.1**

**Technical Reference: Communications, Volume 2**

**Publication No. SC23-6611-02**

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send your comments via e-mail to: [pserinfo@us.ibm.com](mailto:pserinfo@us.ibm.com)

If you would like a response from IBM, please fill in the following information:

\_\_\_\_\_  
Name

\_\_\_\_\_  
Address

\_\_\_\_\_  
Company or Organization

\_\_\_\_\_  
Phone No.

\_\_\_\_\_  
E-mail address



Cut or Fold  
Along Line

Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

**BUSINESS REPLY MAIL**

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation  
Information Development  
Department 04XA-905-6B013  
11501 Burnet Road  
Austin, TX 78758-3400



Fold and Tape

Please do not staple

Fold and Tape

Cut or Fold  
Along Line





Printed in U.S.A.

SC23-6611-02

